

DECEMBER 1983

(12) LIDS-R-1354

AD A138614

PROCEEDINGS OF THE
6TH MIT/ONR WORKSHOP ON C³ SYSTEMS

Reproduced From
Best Available Copy

UHS FILE COPY

EDITED BY:

MICHAEL ATHANS
ELIZABETH R. DUCOT
ALEXANDER H. LEVIS
ROBERT R. TENNEY

20000802053

Laboratory for Information and Decision Systems
MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MASSACHUSETTS 02139

DECEMBER 1983

AD/A138 614
LIDS-R-1354

PROCEEDINGS OF THE
6TH MIT/ONR WORKSHOP ON C³ SYSTEMS

HELD AT
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
CAMBRIDGE, MASSACHUSETTS
JULY 25 TO 29, 1983

EDITED BY

MICHAEL ATHANS
ELIZABETH R. DUCOT
ALEXANDER H. LEVIS
ROBERT R. TENNEY

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

REPRODUCED BY
NATIONAL TECHNICAL
INFORMATION SERVICE
U.S. DEPARTMENT OF COMMERCE
SPRINGFIELD, VA. 22161

SPONSORED BY

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LABORATORY FOR INFORMATION AND DECISION SYSTEMS
CAMBRIDGE, MASSACHUSETTS

WITH SUPPORT FROM

OFFICE OF NAVAL RESEARCH

CONTRACT ONR/N00014-77-C-0532(NR041-519)

This document has been approved
for publication
distribution

DISCLAIMER NOTICE

**THIS DOCUMENT IS BEST QUALITY
PRACTICABLE. THE COPY FURNISHED
TO DTIC CONTAINED A SIGNIFICANT
NUMBER OF PAGES WHICH DO NOT
REPRODUCE LEGIBLY.**

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
	AD-A138 614	
4. TITLE (and Subtitle)	5. TYPE OF REPORT & PERIOD COVERED	
6TH MIT/ONR WORKSHOP ON C ³ SYSTEMS	INTERIM	
	6. PERFORMING ORG. REPORT NUMBER	
	LIDS-R-1354	
7. AUTHOR(s) Editors	8. CONTRACT OR GRANT NUMBER(s)	
Michael Athans Robert R. Tenney Elizabeth R. Ducot Alexander H. Levis	ONR-N00014-77-C-0532	
9. PERFORMING ORGANIZATION NAME AND ADDRESS	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
Laboratory for Information and Decision Systems Massachusetts Institute of Technology Cambridge, Massachusetts 02139	NR-041-519	
11. CONTROLLING OFFICE NAME AND ADDRESS	12. REPORT DATE	
Office of Naval Research Arlington, Virginia 22217	December 1983	
	13. NUMBER OF PAGES	
	xi + 293	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report)	
	UNCLASSIFIED	
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report)		
Approved for public release: Distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
COMMAND, CONTROL, COMMUNICATIONS		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		
This report contains printed manuscripts of papers presented at the Workshop by several authors.		

DD FORM 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED
SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

COMPONENT PART NOTICE

THIS PAPER IS A COMPONENT PART OF THE FOLLOWING COMPILATION REPORT:

(TITLE): Proceedings of the MIT/DNR (Massachusetts Institute of Technology/Office of Naval Research) Workshop on C3 (Command, Control, Communications) Systems (6th)
Held at Cambridge, Massachusetts on 25-29 July 1983.

(SOURCE): Massachusetts Inst. of Tech., Cambridge, Lab. for Information and Decision Systems.

TO ORDER THE COMPLETE COMPILATION REPORT USE AD-A138 614.

THE COMPONENT PART IS PROVIDED HERE TO ALLOW USERS ACCESS TO INDIVIDUALLY AUTHORED SECTIONS OF PROCEEDINGS, ANNALS, SYMPOSIA, ETC. HOWEVER, THE COMPONENT SHOULD BE CONSIDERED WITHIN THE CONTEXT OF THE OVERALL COMPILATION REPORT AND NOT AS A STAND-ALONE TECHNICAL REPORT.

THE FOLLOWING COMPONENT PART NUMBERS COMPRISE THE COMPILATION REPORT:

AD#:	P002 875	TITLE: Doing C2 (Command Control) Experiments Using War Games.
	P002 876	A Comparison of Two Air Defense Command and Control Models.
	P002 377	Lanchester's Equations and Game Theory.
	P002 878	Information Processing and Decisionmaking Organizations: A Mathematical Description.
	P002 879	A Model for Asynchronous Distributed Computation.
	P002 880	A Unified Approach to Modeling and Combining of Evidence through Random Set Theory.
	P002 881	Software for Explicitly Probabilistic Mathematics.
	P002 882	Command Decision Makers and Their Modes of Interaction.
	P002 883	Information Theoretic Models of Memory in Human Decisionmaking Models.
	P002 884	Human Decisionmaking in Dynamic Environments with Increasing Information Processing Demands.
	P002 885	Individual Differences in Military Decisionmaking: The Classification Performance of Active Sonar Operators.
	P002 886	Decision and Display Analysis in a Simple Surveillance Problem.
	P002 887	The Effects of a Spatial Information Format on Decision Making Performance in a C3 (Command, Control, and Communications) Probabilistic Information Integration Task.
	P002 388	Development of a Generalized Human-Machine Interface.
	P002 389	Tactical Decision Aids Which Quantify Judgment.
	P002 890	A Knowledge Based Interactive Procedure for Planning and Decision Support under Uncertainty and Parameter Imprecision.
	P002 891	An Evaluation of ARIADNE (Alternative Ranking Interactive Aid Based on Dominance Structure Information Elicitation).
	P002 892	Research on Cognitive Collaboration between Persons and Computers.
	P002 893	Coordination of Battle Group Warfare Commanders through Summary Display Transfers.
	P002 894	A Man-Machine Interface Concept for a State-of-Art, Shipboard, Command/Control Console.

FOREWORD

The Sixth MIT/ONR C^3 Workshop was held from July 25, 1983 through July 29, 1983. Sessions on the first four days were held at the Massachusetts Institute of Technology, Cambridge, Massachusetts, and these Proceedings constitute the written record of the work presented there. The fifth day's session was held at The MITRE Corporation, Bedford, Massachusetts, and was organized separately by the Office of Naval Research. As the topics under discussion there were classified, the corresponding papers have not been included in this volume.

The Workshop attracted approximately 160 persons; its greatest strength was the variety of backgrounds represented by these participants. Leaders from government laboratories and funding agencies, managers and technical staff of high technology firms, and faculty, staff, and students from academic institutions, all came together to contribute to the steady evolution of C^3 from an art to a science.

Several themes continued to appear throughout this Workshop; themes which can be traced back through the preceding ones. These are reflected in the section headings used in the organization of the proceedings. The most challenging problem facing C^3 scientists is that of devising means to evaluate the performance of alternative C^3 system architectures. This must take into account the sensors, communications, and weapons available to friendly as well as enemy forces. It must involve a closed-loop, cybernetic model of the entire military situation and its must inevitably involve a deep understanding of the behavior of the humans which execute the key decision functions from sonar operators through battle group commanders. From the perspective of design, surveillance has been reduced to a relatively precise science, at least compared with the design of the decision function.

The purpose of these Workshops is to foster real-time interaction between many of the key contributors in the developing science of C^3 . In this respect, the Sixth Workshop was a significant success. We look forward to the Seventh Workshop, to be held from June 11 to June 15, 1984 in sunny San Diego, California, to be equally successful.

Robert R. Tenney
Cambridge, Massachusetts
December 1983

COMPONENT PART NOTICE (CON'T)

AD#: P002 895	TITLE: Multi-Sensor Fusion, Communications and Information Warfare.
P002 896	Techniques for Detecting Cover and Deception.
P002 897	Distributed Estimation in the MIT/LL (Massachusetts Institute of Technology/Lincoln Laboratory) DSN (Distributed Sensor Networks) Test-Bed.
P002 898	Distributed Estimation Systems.
P002 899	Optimal Maneuver Detection and Estimation in Multiobject Tracking.
P002 900	Generalized Tracker/Classifier (GTC) - A System for Tracking and Classification of Multiple Targets by Multiple Sensors.
P002 901	A Framework for Evaluation of Surveillance Algorithms.
P002 902	INCA (Intelligent Correlation Agent): An Environment for Exploratory Development of Tactical Data Fusion Techniques.
P002 903	Track Association Algorithm.
P002 904	Further Results in Multiplatform Correlation and Gridlock in the Naval Battlegroup.
P002 905	Inter-ship Sensor Alignment Using UHF Tactical Data Link Relative Navigation.
P002 906	The Use of Reciprocal Polar Co-Ordinates in Passive Tracking.
P002 907	Stochastic Control of a Passive Search for Ocean Vehicles.
P002 908	Detection Modeling of Unsaturated and Partially Saturated Ocean Acoustic Signals.
P002 909	The Design of Information Structures: Basic Allocation Strategies for Organizations.
P002 910	Informational Aspects of a Class of Subjective Games of Incomplete Information.
P002 911	A Minimum Sensitivity Incentive Control Approach to Team Problems.
P002 912	Dynamic, Hierarchical Decision Problems.
P002 913	Optimal File Allocation Problems for Distributed Data Bases in Unreliable Computer Networks II.

TABLE OF CONTENTS

	<u>Page</u>
FOREWORD.....	v
SECTION I: Evaluation: Methods and Tools	1
1. DOING C ² EXPERIMENTS USING WAR GAMES	3
Joel S. Lawson, Jr. Naval Electronic Systems Command	
2. AUTOMATED WAR GAMING AS A TECHNIQUE FOR EXPLORING STRATEGIC COMMAND AND CONTROL ISSUES	10
Paul K. Davis, Peter J. E. Stan, and Bruce W. Bennett The Rand Corporation	
3. A COMPARISON OF TWO AIR DEFENSE COMMAND AND CONTROL MODELS	18
John R. Dowdle, Leslie C. Kramer, Robert F. Gendron, and Mark P. Merriman ALPHATECH, Inc.	
4. LANCHESTER'S EQUATIONS AND GAME THEORY	24
J. M. Vosenkraft and P. H. Moose Naval Postgraduate School	
5. INFORMATION PROCESSING AND DECISIONMAKING ORGANIZATIONS: A MATHEMATICAL DESCRIPTION	30
Alexander H. Lewis Massachusetts Institute of Technology	
6. A MODEL FOR ASYNCHRONOUS DISTRIBUTED COMPUTATION	39
Dimitri P. Bertsekas Massachusetts Institute of Technology	
7. A UNIFIED APPROACH TO MODELING AND COMBINING OF EVIDENCE THROUGH RANDOM SET THEORY	42
I. R. Goodman Naval Ocean Systems Center	
8. SOFTWARE FOR EXPLICITLY PROBABILISTIC MATHEMATICS	48
Steven M. Goldstein The MITRE Corporation	
SECTION II: Humans: Models and Interfaces	59
9. COMMAND DECISION MAKERS AND THEIR MODELS OF INTERACTION	61
Peter D. Morgan SCICON, Ltd.	

10. INFORMATION THEORETIC MODELS OF MEMORY IN HUMAN
DECISIONMAKING MODELS 67
Susan A. Hall, RCA Astroelectronics
Alexander H. Lewis, Massachusetts Institute of Technology
11. HUMAN DECISIONMAKING IN DYNAMIC ENVIRONMENTS WITH INCREASING
INFORMATION PROCESSING DEMANDS 76
Daniel Serfaty, Eric P. Souleby, and David L. Kleinman
University of Connecticut
12. INDIVIDUAL DIFFERENCES IN MILITARY DECISIONMAKING:
THE CLASSIFICATION PERFORMANCE OF ACTIVE SONAR OPERATORS 83
Joseph G. Wohl
ALPHATECH, Inc.
- 1135857 13. THE COGNITIVE ORGANIZATION OF SUBMARINE SONAR INFORMATION:
A MULTIDIMENSIONAL SCALING ANALYSIS 87
Kevin Laxer, George Moeller, and William H. Rogers
Naval Submarine Medical Research Laboratory
14. DECISION AND DISPLAY ANALYSIS IN A SIMPLE SURVEILLANCE PROBLEM 91
Frank L. Greitzer and Ramon L. Horshman
Navy Personnel Research and Development Center
15. THE EFFECTS OF A SPATIAL INFORMATION FORMAT ON DECISION MAKING
PERFORMANCE IN A C³ PROBABILISTIC INFORMATION INTEGRATION TASK 96
Captain Brad Scott, U. S. Military Academy
Christopher D. Wickens, University of Illinois
16. DEVELOPMENT OF A GENERALIZED HUMAN-MACHINE INTERFACE 100
Rita E. Knox
Lockheed Electronics Company, Inc.
17. TACTICAL DECISION AIDS WHICH QUANTIFY JUDGEMENT 106
Rex V. Brown, and Marvin S. Cohen
Decision Science Consortium, Inc.
18. A KNOWLEDGE BASED INTERACTIVE PROCEDURE FOR PLANNING AND DECISION
SUPPORT UNDER UNCERTAINTY AND PARAMETER IMPRECISION 109
Andrew P. Sage and Chelsea C. White, III
University of Virginia
19. AN EVALUATION OF ARIADNE 117
Chelsea C. White, III, Andrew P. Sage,
Shigeru Dozono, and William T. Scherer
University of Virginia
20. RESEARCH ON COGNITIVE COLLABORATION BETWEEN
PERSONS AND COMPUTERS 122
Marvin S. Cohen
Decision Science Consortium, Inc.

21.	COORDINATION OF BATTLE GROUP WARFARE COMMANDERS THROUGH SUMMARY DISPLAY TRANSFERS	129
	Conrad J. Grant The Johns Hopkins University	
22.	A MAN-MACHINE INTERFACE CONCEPT FOR A STATE-OF-THE-ART, SHIPBOARD, COMMAND/CONTROL CONSOLE	133
	Glenn A. Oaga Naval Ocean Systems Center	
SECTION III: Surveillance: Issues and Algorithms		137
23.	MULTI-SENSOR FUSION, COMMUNICATIONS AND INFORMATION WARFARE	139
	Daniel Schutner The Pentagon	
24.	TECHNIQUES FOR DETECTING COVER AND DECEPTION	145
	Robert P. Goldsmith and Ralph F. Gorenz Betac Corporation	
25.	DISTRIBUTED ESTIMATION IN THE MIT/LL DSN TEST-BED	152
	J. R. Delaney, R. T. Lacosz, and P. E. Green Massachusetts Institute of Technology/Lincoln Laboratory	
26.	DISTRIBUTED ESTIMATION SYSTEMS	158
	C. Y. Chong, S. Mori, Advanced Information and Decision Systems Edison Tse, Stanford University	
27.	OPTIMAL MANEUVER DETECTION AND ESTIMATION IN MULTIOBJECT TRACKING	164
	Thomas Karion, Alan Blitz, Robert Washburn, ALPHATECH, Inc. Alan Willsky, Massachusetts Institute of Technology	
28.	GENERALIZED TRACKER/CLASSIFIER (GTC) - A SYSTEM FOR TRACKING AND CLASSIFICATION OF MULTIPLE TARGETS BY MULTIPLE SENSORS	172
	Shozo Mori, Chee-Yee Chong, and Richard P. Wishner Advanced Information and Decision Systems Edison Tse, Stanford University	
29.	A FRAMEWORK FOR EVALUATION OF SURVEILLANCE ALGORITHMS	178
	Jerold L. Weiss, ALPHATECH, Inc. Robert R. Tenney, Massachusetts Institute of Technology	
30.	INCA: AN ENVIRONMENT FOR EXPLORATORY DEVELOPMENT OF TACTICAL DATA FUSION TECHNIQUES	183
	Leonard S. Gross, and Harold J. Payne VERAC, Inc.	
31.	TRACK ASSOCIATION ALGORITHM	188
	A. C. Bamford, E. M. L. Beale, and D. East SCICON, Ltd.	

32. FURTHER RESULTS IN MULTIPLATFORM CORRELATION AND CRIDLOCK
IN THE NAVAL BATTLEGROUP191
M. A. Kovsich
COMPTON Research, Inc.
33. INTERSHIP SENSOR ALIGNMENT USING UHF TACTICAL DATA LINK
RELATIVE NAVIGATION195
R. H. Overton
COMPTON, Research, Inc.
34. THE USE OF RECIPROCAL POLAR CO-ORDINATES IN PASSIVE TRACKING198
A. C. Bamford, E. M. L. Beale, J. Lee, and S. Patel
SCICON, Ltd.
35. STOCHASTIC CONTROL OF A PASSIVE SEARCH FOR OCEAN VEHICLES202
LCDr. G. Z. Wilhoit, U. S. Navy
Robert R. Tenney, Massachusetts Institute of Technology
36. DETECTION MODELING OF UNSATURATED AND PARTIALLY SATURATED OCEAN
ACOUSTIC SIGNALS207
Anastassios N. Perakis, University of Michigan
Marilao N. Psaraffis, and Henry I. Gonzalez
Massachusetts Institute of Technology

SECTION IV: Decision: Organization and Theory213

- 1137476 37. DISTRIBUTED OPTIMIZATION ALGORITHMS WITH COMMUNICATIONS215
John Tsitsiklis, Dimitri P. Bertsekas, and Michael Athans
Massachusetts Institute of Technology
38. THE DESIGN OF INFORMATION STRUCTURES:
BASIC ALLOCATION STRATEGIES FOR ORGANIZATIONS219
Debra A. Stille, ICF, Inc.
Alexander M. Lewis, Massachusetts Institute of Technology
39. INFORMATION ASPECTS OF A CLASS OF SUBJECTIVE GAMES
OF INCOMPLETE INFORMATION226
Demosthenis Teneketzis and David C. Castanon
ALPHATECH, Inc.
- 139 443 40. ON THE COMPLEXITY OF DISTRIBUTED DECISION PROBLEMS232
John Tsitsiklis and Michael Athans
Massachusetts Institute of Technology
- 124 940 41. OPTIMAL DECENTRALIZED CONTROL OF FINITE NONDETERMINISTIC
SYSTEMS238
Robert R. Tenney
Massachusetts Institute of Technology

42. A MINIMUM SENSITIVITY INCENTIVE CONTROL APPROACH TO TEAM PROBLEMS	243
---	-----

Tamer Basar, Jose B. Cruz, Jr., and Derya Cansever
University of Illinois

43. DYNAMIC, HIERARCHICAL DECISION PROBLEMS	249
---	-----

Peter B. Luh, and Taikang Ning, University of Connecticut
Jia-Shuan Chang, SUNY at Stony Brook

44. OPTIMAL FILE ALLOCATION PROBLEMS FOR DISTRIBUTED DATA BASES IN UNRELIABLE COMPUTER NETWORKS II	255
--	-----

Moses Hsingwen Ma and Michael Athans
Massachusetts Institute of Technology

45. A COLLISION RESOLUTION PROTOCOL WITH LIMITED CHANNEL SENSING-FINITELY MANY USERS	268
--	-----

P. Papantoni-Kazikos, Glenn E. Marcus, and Michael Georgiopoulos
University of Connecticut

LIST OF ATTENDEES	273
-------------------------	-----

C ³ WORKSHOP PROGRAM	283
---------------------------------------	-----

AUTHOR INDEX	293
--------------------	-----

SECTION I

Evaluation: Methods and Tools

DOING C² EXPERIMENTS USING WAR GAMES

Dr. Joel S. Lawson, Jr.

Technical Director, C³I Systems & Technology Directorate (ELEX 06T)Naval Electronic Systems Command
Washington, DC 20363INTRODUCTION

In the last few years considerable progress has been made in the development of an analytic theory of military Command Control (C²), both as a process and as a large-scale system. There are now models of C² organizations which permit the examination of the effects of various changes in a C² system and which can predict some of the behavior of such a system in a gross sense. And, due to the increased attention being given the field, with the attendant increase in papers, workshops, etcetera, there is slowly developing a common vocabulary for use in the emerging "C² Theory."

What is still lacking, however, is a body of experimental data which can be used as a "touchstone" to guide further theoretical developments, and against which theoretical predictions can be tested.

This paper reports the results of a very rudimentary experiment which was conducted at the Naval Post-graduate School at Monterey, California, during the 1983 Winter Quarter to test two specific hypotheses. As is often the case, it was found necessary to modify or restate the hypotheses during the conduct of the experiment in order to accommodate certain "real world" constraints. The results, however, are both interesting in their own right and reassuring for the prospect of being able to do further experiments in the C² arena.

Motivation for the Experiments

The practical motivation for this experiment lay in a desire to study the "decision making" process in a C² process. "Decisions" are one of the two major products of a C² system. (The other major C² system product may be considered to be "plans" which are statements of actions to be taken by various parts of the system. They are generated either to implement a previous decision or to provide for possible future decisions. In many cases they have the nature of "pre-determined decisions.") It was this motivation which led to the approach and the experimental procedure adopted.

In keeping with a "hard science" approach to the development of a C² Theory, we wanted to find "observables" which could be observed by people who were not part of the system and which could also be observed by other investigators, in other experiments.

It should be noted that, by the nature of the C² business, which involves both people and complex scenarios, it is nearly impossible to "replicate" an experiment. People, unfortunately, learn, and their behavior on a second trial will be different from what it was on the first. Also, different "command teams" or battle staffs will have different approaches or "styles" in dealing with the same basic problem. Furthermore, in "free play" exercises, which are the only ones in which real decision making takes place, the dynamics of the entities in the scenario very seldom repeat. In particular, successive plays of the same scenario generally do not lead to the same physical configuration of the battlefield after even a few minutes of play. Therefore, "replication," in the sense it is used in physics or chemistry, is a practical impossibility. Rather, we must search for observables which are in some way

connected to the decision-making process and which can be observed in numerous settings, but which are independent of the details of the scenario and the particular participants.

Pursuing this line of reasoning, it was decided that an interesting characteristic quantity might be the time interval between decisions, independent of what the decisions were. A histogram of this interval against the number of cases in which it was observed should give some idea about the tempo of activity in a command center.

But this leads to the question of defining "a decision" and how it is to be observed. At first it was felt that decisions could be grouped into broad classes, such as tactical, strategic, or logistics, and observed or logged by listening to the conversations between members of the battle staff. In practice, this turned out to be virtually impossible, and a simpler but more exact definition of a decision was adopted.

Specifically, decisions were defined as space-time events made manifest by either utterances or a series of keystrokes on a computer terminal which request or transmit information or directions. This definition satisfies the requirement that the decision be observable, i.e., it is given physical reality. (My decision to have steak for dinner tomorrow has no reality until I make it known to the outside world by telling someone about it.) This choice also takes us down to the "operational level" in the decision-making process. A decision to ask for information, or to send a message, is now treated equally with a decision to engage the enemy. (Perhaps a better name for "decisions" defined in this way would be "transactions," but for now we shall continue to call them decisions.) This definition of decision led directly to a revised version of the original first hypothesis:

Preceding page blank

A Commander or Battle Staff produces decisions which are evidenced by utterances or keystrokes and which can be observed by external observers. Furthermore, different external observers will agree on the "class" to which the observed decision belongs if the classes are chosen in a reasonably general manner.

The notion of a "class" of decisions will be described further below, when we discuss the experimental procedure. But before going on to that, it will be recalled that there was a second hypothesis to be tested in this experiment. The original form of this hypothesis was stated as follows:

The histogram of time intervals between decisions should give an indication of the state of training or competency of the Commander or Battle Staff.

It seemed reasonable to assume that a relatively untrained team would take almost random amounts of time to arrive at each of a series of decisions, leading to a nearly flat histogram. At the same time, a more experienced group would have better defined work habits and evidence a more peaked distribution. In fact, one might expect a bimodal distribution, with simple or routine decisions being made quickly and effortlessly, while more complex procedures took longer. Alternatively, a complex problem might result in a long time delay, followed by a flurry of activity with short apparent inter-transaction times.

Unfortunately the data collected in this experiment do not allow for a simple determination of which decisions may have resulted from some previous one. Nor is there any direct way to associate one decision with another in an input-output sense. Investigation of these matters will have to await an improved experimental procedure.

The Experimental Setting

An opportunity to test these hypotheses arose during the 1983 Winter Quarter at the Naval Postgraduate School, Monterey, California.

A series of computer-based war games were to be run, using the Warfare Environment Simulator (WES) at the Naval Ocean Systems Center (NOSC), San Diego, California, to which the Postgraduate School has a remote access capability. The purpose of these games was to give the students in the Operations Research curriculum an opportunity to get "hands-on" experience with the operation of a war game. Thus, the subject "Battle Staffs" were certainly inexperienced and, in fact, being made up of members of all four services, many of the players were unfamiliar with Navy terminology and procedures.

The WES facility at the Postgraduate School at that time consisted of three sets of displays, one each for the Orange and Blue teams and one for the umpire or game coordinator. Each set of displays had a large graphic display of the theater of operations in the center with a smaller alphanumeric display and keyboard on each side of it. The scale of the geographic display could be changed to provide a "zoom" capability and appropriate symbology depicted the location of objects. One of the alphanumeric displays was generally used only as a status board, while the other was used to give instructions to the machine to effect control of that side's forces.

The display systems in the Blue and Orange command posts showed only those targets which had been detected by their own sensor systems, while the umpire's position had access to the data as seen by each side as well as to ground truth.

The class running the war game was split into three groups (A, B, and C) of two (Orange and Blue) three-man teams. Each was a familiarization session in which they learned how to operate the keyboard and make the machine do some elementary things. In the second session they actually did play a war game, deploy their forces, and get to the point of engaging the enemy. The final session was run "for the record" and was based on a scenario and initial force disposition very similar to the ones with which they had practiced. By this time the Battle Staffs seemed to be reasonably comfortable with the system and could concentrate most of their effort on the battle rather than the game mechanics.

To carry out the experiment and test these hypotheses, observers were recruited from other students enrolled in the Command Control and Communications curriculum at the Postgraduate School, and arrangements made to have them present during the running of the games.

Only limited observations were made during the first series of runs, to test the methodology. The data reported here were all taken during the second and third series of runs, when the performance of the Battle Teams had become reasonably stable.

Experimental Procedure

For the experiment, each observer was equipped with a clipboard of data sheets which had a synchronized clock mounted on it. They were instructed to record to the nearest second when they thought they observed a "decision" and the "class" of that decision. A sample of the data sheet is reproduced in Figure 1, which also shows the classes of decisions used in this experiment. These classes were intended to be generic in nature and encompass nearly all the situations which were expected to be observed.

The observer was also asked, if he had time, to enter any amplifying information in the "Notes" column. The columns for "Game Time" and "Enter Time" were provided to allow correlation to game time (which can be faster or slower than clock time) and to note the time when an instruction which had been observed passing from one team member to another, was actually entered on the keyboard. The latter turned out to be an impossible task as the keyboards were in such constant use as to preclude connecting any particular set of keystrokes with any specific utterance. (In future experiments it may be possible to achieve some of these measurements by installing appropriate "hooks" in the game software.)

After the experimental session, the observations were transcribed into a computer program which computed the desired time intervals and could generate various histograms of the results. Figure 2 is a sample of a data file generated by this program.

As shown in the Figure, it was eventually decided to classify the decisions by "type" as well as "class." This was partly due to the confusion between different observers as to which class a particular decision should be placed in, and partly because of the limited number of observations of members of a particular class. By introducing a definition of "types" of decisions as either "information

decisions" or "action decisions," the sample size was effectively increased and the differences between observers was reduced to a negligible amount. Examples of this distinction are decisions which deal with "information," such as requesting status or identity, and those decisions which deal with "action," such as deploying a force or changing the ENCOM condition.

With these additions to the data definitions, the time intervals are defined as follows:

- T0 = real world clock time
- T1 = time since last decision, independent of its class
- T2 = change in T1 since last T1 event
- T3 = time since last decision of same "type"
- T4 = change in T3 since last T3 event
- T5 = time since last decision of same "class"
- T6 = change in T5 since last T5 event

By generating such data files on disk, each representing the observations of one observer (using the observer's name as file name) during a game he observed (using the team name and game number as a file extension), the data was set up to allow comparison of different observers as well as different teams and games.

The results of this analysis are reported in the next section.

Experimental Results

Before examining the specific experimental results, a word about the choice of "classes of transactions" is in order because they turned out to influence the apparent results. In particular, the distinction between ordering the deployment of one or two units or platforms and ordering the deployment (e.g., a course change) of the whole force was interpreted differently by different observers. And the same was true of requests for the status of individual units or of a major portion of the force. Much of this confusion could have been overcome if there had been more time to test the methodology and have the data takers agree on some conventions. As it was, there was a minimum amount of coordination between the observers, and each one made up his own rules as to how he would classify the events he observed.

The basic classes were originally chosen because it was assumed that they were sufficiently general to encompass all the transactions which were likely to take place, while also being representative of distinctly different sorts of activity which one would expect to see in a command center.

In practice, it appears that these were reasonable assumptions, but that the definition of the "classes of transactions" can probably be improved, and in fact should probably be tailored to match the nature of the scenario and/or level of command which is being examined. Figures 3 and 4 present a typical comparison of the time between transactions, independent of the type or class of transaction, as recorded by two different observers watching the same battle staff. These Figures are histograms of the time intervals between observed decisions, grouped into 18-second wide bins. (The units of "width" are seconds.) For ease of comparison, the actual number of observations has been normalized to 100 so these represent percentage distributions.

It is immediately obvious that the observers were apparently observing the same sort of behavior and

recording a similar series of events. A detailed comparison of their event logs bears this out. In fact, in all cases where there were two or more observers watching the same battle staff, there was remarkable agreement in what they logged as "decisions" or transactions.

The two major sources of nonconformity of the logs were the differences of interpretation mentioned above and a very simple but important physical limitation. Because the battle staff were sitting side-by-side in front of the displays, the observers tended to sit on either side of them so they could both see the displays and hear what the group was talking about. This led to one observer being able to hear remarks about what was on the status board better than the other, while the second was more alert to directions to the keyboard operator.

Nonetheless, as a general rule the overlap of the logs of pairs of observers was about 60 percent, even without any serious attempt to agree on just how the classification scheme would be applied. At the grosser level of whether it was an "action" transaction or one dealing with information, the agreement was nearly 80 percent. The "disagreements" are nearly all accounted for by a missing entry on one of the logs. That is, one observer logged an event which the other did not. In most cases this can be explained as a result of the placement of the observer relative to the battle staff, as mentioned above.

Moreover, data taken by the same observer watching different Battle Staffs produces plots very similar to those shown in Figures 3 and 4. So the distribution of time intervals between decisions does seem to be a characteristic quantity, at least in this setting.

An interesting change is observed, however, when the time between successive action or information transactions is plotted, as in Figures 5 and 6. The action decisions still show the Rayleigh-like distribution, while the information decisions are more uniformly distributed. No explanation for this is offered, but an interesting comparison with some other data will be made below. (Also it was noted that some observers seemed to be more sensitive to action decisions than to information decisions, perhaps reflecting their personality more than their location in the command space.)

Based on these results, it seems that the first of the hypotheses to be tested is confirmed by the data which was taken during this experiment.

As far as the second hypothesis is concerned, we can only report that there is no significant difference between the data obtained in the second and third sets of war games. If there was any "learning effect," it was masked by other attributes of the game.

A final interesting observation is that in essentially all cases the observers logged about two-thirds of the transactions as action decisions and only one-third as information decisions. This may be partly due to the nature of the WES facility, in which the Battle Staffs have to issue a rather large number of maneuvering instructions to their (computer-simulated) forces. This has the effect of putting the Battle Staff in a much more "tactical" role than a "command" role. Also, with only one command center on each side, there were no superiors or subordinates to answer questions for or receive

queries from. A third cause may have been the presence of electronic status boards which provided information without the verbal clues which permitted the transaction to be recorded.

Comparison With a Manual War Game

Coincident with the conduct of this experiment, the "historical logs" of a game run at the war-gaming facility at the Naval War College at Newport, Rhode Island, during the preceding fall became available.

This game also involved some inexperienced players but these were supported by trained and experienced Naval officers, and both sides had had considerable time to develop their plans before the game commenced. In addition to the game being basically a manual one without automated status boards, etc., there were several levels of command represented on each side. Thus it was quite a different environment and perhaps more closely simulated the conditions which one might expect in a typical command node in a Navy C2 system.

In order to obtain some comparative data, the logged events were divided into the classes shown in Figure 7. These classes were chosen as being more representative of the log entries, although they still retain the distinction between action and information decisions. (Although these classes are also divided into "input" and "output" groups, no analysis of this dimension has yet been undertaken.)

Using these data, which are based on the editing and compilation of many observers' logs into one by the game historian, similar analyses were carried out. Two typical plots are shown in Figures 8 and 9, which depict one six-hour period in the game at two levels of command, CTF 70.1, the Battle Force Commander, and 70.X, one of his four immediate subordinates. It should be noted that the basic time unit of "width" used in these Figures is one minute rather than the second used for the WES data.

While the most striking thing about these data is the similarity of the plots to those obtained with the data taken during the WES war games, there are several other interesting points.

First, the activity seems to be peaked at the level of the Battle Force commander. While not displayed, the plots of activity at the level of SEVENTH Fleet Commander show only about one-sixth the number of decisions per unit time that are evident at the Force Commander level. And the Force Commander's subordinate is operating at about half his tempo. However, there are four subordinates, so the total activity is higher at the lower level, as one would expect it to be.

Second, the effective mean time between decisions seems to be about two minutes, rather than the 15 to 20 seconds seen in the WES game. This, of course, may be an artifact of the coarser time increments used in the Newport logs, as well as the less tactical nature of the decisions being made.

And thirdly, the statistics on types of transactions were exactly reversed from that seen at WES. Two-thirds of the decisions involved information and only one-third or less involved action. Whether this was due to the lack of computer aids or to a more realistic simulation of a command situation is not clear. It is hoped that additional light can be shed on this subject by obtaining some more data at the War College, which has had a major computer

facility installed to support its war gaming functions since these data were taken.

Conclusions

First, it seems evident that one can define space-time events (which we shall call decisions or transactions) which deal with either actions or information and which are made observable through utterances or keystrokes made by the Battle Staff, on which different observers will agree. That is, different observers will report observing the same type of event at the same time. And it seems likely that these space-time events could be further subdivided with more careful training and discussion among the observers.

Second, the gross inter-transaction time has a distribution not unlike a Rayleigh distribution, and which seems to be a common behavior even in quite different situations.

Third, the inter-transaction time between transactions of the same type may differ between different types.

Fourth, the time between transactions seems to be a minimum at the level of the highest "on-scene" Commander. This is interpreted as meaning that his "operating tempo" is forced to be high enough to encompass the total activity below him, which is divided among several subordinates.

And finally, the qualitative differences between the Newport data and that taken on WES may indicate that introducing computer assistance into a command center will change the nature of the "transactions" which one sees taking place. If true, this may have serious implications for the design and organization of future command control systems.

Group	Side	Date	Observer	Page
Decision Classes				
ES = request EMCON status		CE = order a change in EMCON		
FS = request Force status		DF = deploy the Force		
US = request status of a unit		DU = deploy a unit		
RI = request for an identity		EE = engage enemy (or target)		
SD = seek other data		AO = all other decisions		
Decision Time	Decision Type	Notes	Entry Time	Game Time

Figure 1.
Data Form Used for WES Experiments.

COMMAND CENTER INTER-ACTIVITY TIMES										
No.	DTG	T0	T1	T2	TYP	T3	T4	Act	T5	T6
0	151800	0	0	0	0	0	0	0	0	0
1	151833	33	33	33	2	0	0	AO	0	0
2	152317	317	284	251	1	0	0	FS	0	0
3	152406	366	49	-235	2	333	333	AO	333	333
4	152432	372	26	-23	1	75	75	FS	75	75
5	152508	428	36	10	1	36	-39	FS	36	-39
6	152530	450	22	-14	2	84	-29	DU	0	0
7	152612	472	42	20	1	64	28	SD	0	0
8	152735	575	83	41	1	83	19	US	0	0
9	152807	607	32	-51	1	32	-51	SD	115	115
10	152854	654	47	15	1	47	15	US	79	79
11	152937	697	43	-4	2	247	163	AO	331	-2
12	153015	735	38	-5	2	38	-209	AO	38	-293
13	153058	778	43	5	2	43	5	DF	0	0
14	153146	826	48	5	2	48	5	DF	48	48
15	153256	896	70	22	2	70	22	AO	161	123
16	153323	923	27	-43	1	269	0	FS	495	0

Figure 2.
Print-Out of Computed Time Intervals as Stored in Computer Data Files.

```

0      10      20      30      40      50
I....I....I....I....I....I....I....I....I....I
I*****
I*****
I*****
I*****
I***
I*****
I***
I*
I
I***
I
I
I*
129 Observations 116 Samples 12 Bins      Width = 18
File = WREN/B2      Act.type = 1      Time int. = T1
Game interval from: 153000 to: 170500      NORMALIZED

```

Figure 3.
Distribution of Time Intervals Seen by Observer Wren.

	Input transactions	Output transactions
Info:	RR = receive a report RQ = receive a query	SR = send a report SQ = send a query
Actions:	RD = receive a directive RT = receive tasking	SD = send a directive ST = send tasking
Miscel:	AD = all other decisions	JE = Journal entry

Figure 7.
Decision Taxonomy Used for Newport data.

The distinction between directives and tasking is that directives leave the method of implementation to the subordinate, while a tasking includes specific detailed instructions. The JE category was used for remarks by the historian.

```

0      10      20      30      40      50
I....I....I....I....I....I....I....I....I....I
I*****
I*****
I*****
I*****
I*****
I***
I**
I
I*
I
I*
I**
174 Observations 174 Samples 12 Bins      Width = 1
File = CTF781/P1  Act.type = 1      Time int. = T 1

```

Figure 8.
Distribution of Time Intervals Observed in
Command Center of Battle Force Commander.

```

0      10      20      30      40      50
I....I....I....I....I....I....I....I....I....I
I*****
I*****
I*****
I*****
I*****
I*****
I***
I**
I*
I*
I*
I**
I*****
89 Observations 89 Samples 12 Bins      Width = 1
File = CTF783/P1  Act.type = 1      Time int. = T 1

```

Figure 9.
Distribution of Time Intervals Observed in the
Command Center of a Subordinate Battle Group Commander.

A
P
2

AUTOMATED WAR GAMING AS A TECHNIQUE
FOR EXPLORING STRATEGIC COMMAND AND CONTROL ISSUES

Paul K. Davis, Peter J.E. Stan, and Bruce W. Bennett

The Rand Corporation
1700 Main Street
Santa Monica, California 90406

Summary

This paper describes a preliminary concept for exploring strategic command and control issues with the automated war gaming of Rand's Strategy Assessment Center (RSAC). The concept features: a top-down functionally oriented approach relevant to the interests of civilian and military leaders; a hierarchical and otherwise multilevel gaming structure; and heuristic rule-based models using a variety of artificial intelligence techniques. The approach will be sensitive to key features of war plans and control procedures. It will make a start on reflecting such phenomena as nonunitary decisionmaking, deception, and confusion. It will take into account some of the asymmetries distinguishing the U.S. and Soviet approaches to C³I. Initial versions of the implemented concept should be useful and interesting but will be relatively simple; with time, it should be possible to evolve gracefully and use some of the detailed models available on pieces of the overall C³I problem. Even the early work, however, will represent a major break with past strategic analysis in which C³I issues have been largely ignored but for limited treatment of communications.

I. Introduction

This paper describes a concept for using the emerging technique of automated war gaming [1] to explore problems of strategic command, control, communications, intelligence, and warning (abbreviated here as C³I or as "command and control"). The concern here is largely with the architecture of an approach rather than with details, many of which remain to be worked out. Formal architecture is essential because we seek an approach that will be broad and useful in its earliest manifestations and that can evolve smoothly over time to address a substantial fraction of the strategic command-and-control issues of principal concern to national leaders. These include: continuity of government; timely command decisions involving both intercontinental and theater-nuclear forces; continuing control over those forces; and prosecution of conflict (which, depending on circumstances, might call for decisive military action, controlled escalation, or de-escalation). In future work we plan to go beyond the current paper's emphasis on strategic nuclear weapons and to extend the architecture to cover strategic aspects of global conflict generally.

The paper's outline is as follows: first, we describe succinctly the principal features of the RSAC automated gaming system; we then discuss our view of what the strategic C³I problem really is--or should be considered to be in our work. Finally, we describe the philosophy of our approach and sketch out our intended plan for implementing it.

II. Background

The Rand Strategy Assessment Center (RSAC) is developing a new approach to strategic analysis that attempts to combine the contextual richness and operational complexity of war gaming with the rigor and transparency of analytic modeling. On the one hand, we are building a large-scale simulation model with the structure of a political-military war game. In this simulation, models represent the various national players, making decisions of both a political and military nature. The simulation can be fully automated. On the other hand, much of the RSAC's work will be highly interactive with human teams playing against computerized adversaries or changing assumptions about such matters as combat outcomes to see the strategic reaction of the automated players.

At the technical level, the RSAC is extending several modern techniques in artificial intelligence (A.I.) as well as using more standard modeling and analysis techniques. [1,2] We shall not discuss the techniques here. Instead, let it suffice to say that: (1) we make extensive use of heuristic rule-based modeling in an English-like programming language; (2) that our decision models use such devices as pattern matching and search (with lookaheads accounting for likely opponent behavior); and, very importantly, (3) that contact with military realism is achieved in part by relating a (greatly extended) version of A.I. scripts to analytic constructs akin to war plans. [1-3] From the viewpoint of A.I. research our effort is notable because it is a rare application to realistic high-level military issues and because the application's scope has caused us to develop concepts for managing complexity in rule-based models that should have more general value. [4]

Figure 1 provides an overview of the RSAC system emphasizing its hierarchical structure, something that will prove important in treating C³I. The first column shows the nominal move sequence in the overall game. The automated players are: (1) Red Agent representing the Soviet Union; (2) Blue Agent representing the United States; (3) Scenario Agent representing to first order all nonsuperpower countries on a country-by-country basis; and (4) Force Agent. The latter model is not really a player--rather, it keeps track of forces worldwide and computes the results of battle and other military operations such as movement. Its individual submodels are typically simple and aggregated, but because of the simulation's breadth and the requirements to interrelate phenomena across theaters, force types, and levels of conflict, Force Agent is quite complicated overall. There is also another model called System Monitor, which guides game development by scheduling moves and managing interfaces with automated recordkeeping, displays, and humans.

The second column provides a closer view of what happens in a single major-agent (Red or Blue) move.

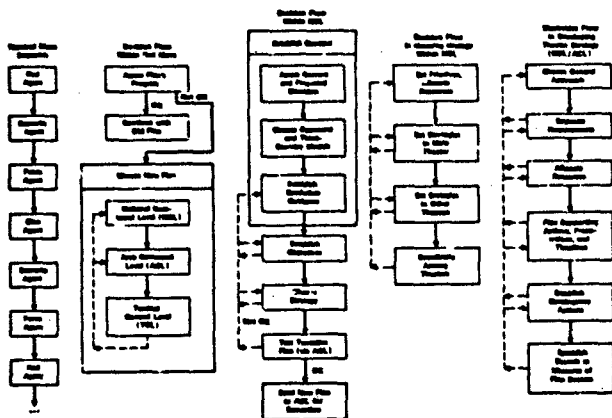


Fig. 1--A Hierarchical View of RSAC Automated War Gaming

The move begins with the agent assessing his success with a previously chosen plan. If all is going well, he merely continues on that plan--which is represented in code by RSAC extensions of A.I. scripts.[2] All plans (or scripts) have bounds, however, and if any of the bounds have been broken (e.g., by excessive attrition or delay, or by the opponent's escalation), then the Agent must reconsider.* This process begins with a rule set associated with functions of the national command level (NCL). The NCL chooses a tentative and incomplete war plan to be filled out and tested by the area command level (ACL), which corresponds loosely with the functions of area commanders such as U.S. CINCs or Soviet TVD commanders. The plan testing includes a lookahead implemented through the tactical control level (TCL), which controls the interfaces with Force Agent and (together with Force Agent's submodels) determines many of the detailed decisions about orders of battle, allocation of resources, etc. (decisions that should not be highlighted in a strategic level game). The lookahead is a game within a game using the agent's assumptions about other players' actions and the likely results of combat. If the plan passes the test, it is then implemented, again through the TCL level. Otherwise, the ACL may adjust the plan and try again, or report back to the NCL that some strategic-level decisions must be changed.**

Continuing to unpeel the onion, the third column provides more detail on what happens in the NCL. The notable feature here is the process model guiding the structure of rule sets. In practice, we must rigidly define the permitted forms of escalation guidance, objectives, and strategies, and then write unambiguous rules leading from game observables (e.g., combatants, location of conflict, and status of forces) to unique permitted forms. We shall not discuss such matters here even though they are consuming a major amount of effort and time.

The fourth column expands upon the Choose Strategy process by noting that strategies must be chosen for each of the several military theaters, including the intercontinental and space theaters, and then

* The techniques for building such plans and scripts are under development by William Jones, Norman Shapiro, and Richard Wise.

** The sophisticated reader will recognize that Fig. 1 is an idealization with imperfect fidelity to the actual computer programs. In the code, distinctions among NCL, ACL, TCL, and decision rules in Force Agent are sometimes fuzzy and the functions alluded to in the boxes of Fig. 1 are sometimes accomplished by rules distributed throughout the program.

coordinated. Finally, the last column expands upon this by suggesting some of the many steps required to define theater-strategy components that would be decided (or at least reviewed) at the NCL. These would include: (1) consistency of actions with overall escalation guidance and objectives; (2) cross-theater coordination; and (3) resource allocation across theaters.

Given this quick overview of RSAC system architecture, let us note some particular items relevant to what follows:

- o Variable behavior patterns. The behavior patterns of Red, Blue, and Scenario Agents are variable to reflect fundamental uncertainties about the true patterns to be expected. Hence, we speak of alternative "Ivans," "Sams," and third-country "temperaments."
- o Parametric force models. Similarly, Force Agent's component models are highly parametric with the parameters chosen for strategic-level analysis (e.g., a few simple equations that calculate bomber prelaunch survivability rather than a complex model considering details such as the propulsor characteristics of a Soviet SLBM that might be used to attack bomber bases).
- o Use of scripts. The decision models do not generally extend below important operational-level issues. Instead, the Agents choose among discrete war plans in the form of scripts, each of which contains a number of microscopic (and sometimes arbitrary) action instructions (which may contain slots for parameters to be filled in by some ACL or Force-Agent subroutine at the appropriate time).
- o Unmodelable Phenomena. RSAC games allow certain phenomena to occur by fiat if the analyst so chooses--e.g., in some fraction of game runs the analyst may want to have riots occur in Poland if certain other events occur. Although the origin and nature of those riots are not simulated, Red's rules may be sensitive to whether the riots occur. The riot flag is, of course, a surrogate for whole classes of important real-world events contributing to fog-of-war effects (and escalation).* Such devices are familiar in manual gaming but may seem unnatural to traditional modelers.

Finally, let us summarize the RSAC system's essential elements, distinguishing between variables, hard-wired items, etc. This is actually not so simple because the RSAC system is designed for flexibility with a variety of users who have different notions about what should be hard-wired. However, Table 1 shows an illustrative breakdown appropriate for an application comparing nuclear-employment concepts in a range of scenarios. Although only illustrative, it demonstrates that there are many pieces to the overall system--all of which must be considered when attempting to treat command and control.

III. Defining the Strategic C³I Problem

Before deciding on an approach to the C³I problem we must first know what it is. One way to define strategic C³I is to list the ingredients of a C³I system: (1) decisionmaking bodies and their procedures; (2) command centers to integrate information for the decisionmakers; (3) control procedures to assure that

*See Refs. 5 and 6 for the RSAC's conceptual approach to escalation modeling.

Table 1

ELEMENTS OF THE RSAC SYSTEM AS VIEWED WHEN ASSESSING
ALTERNATIVE STRATEGIES

Fixed Structure	Fixed Characteristics[a]	Variable Characteristics	Principal Variables
o Game structure built around two (not N) primary players	o War plans/Scripts	o Sams, Ivans, national temperaments	o Employment strategies
o Rules for determining move sequence	o Force models	o Initiating Scenario	
o Sequence of some NCL decisions	o Individual Scenario-Agent rules	o Some key parameters and rules in Force, Scenario, and Scripts	
o Treatment of command levels			

[a] In other applications these would be considered variable.

decisions will be implemented if transmitted; (4) intelligence systems to provide strategic and tactical warning of attacks and other information on enemy forces; and (5) communications to provide information on one's own forces, to permit transmission of decisions, and to permit report-back on the results of execution.

It is important to remind ourselves that C³I does not depend on physical communication systems alone--it involves analysis by the NCA and his staff, as well as lower command levels, decisionmaking procedures, doctrine (which is a key element in control and planning), and many other items. Indeed, the breadth and complexity of command and control are such that definitions and flow diagrams are often too abstract to communicate what the down-to-earth problems people worry about (or should worry about) actually are. It is therefore useful to list some of those illustrative concerns, drawing on the public literature as well as our own knowledge.* In listing these questions, we have not attempted to order them by their actual or perceived importance. Further, although these questions have caught the common fancy as "important" examples of potential C³I failures, we cannot exclude the possibility that prospective problems are more imagined than real in some cases.

Illustrative Questions

1. What if the Soviets attack Washington on Inauguration Day? Would we be "decapitated"? What does "decapitation" mean on an operational level?
2. What if the Secretary of Education becomes the National Command Authority (NCA)? Will he know enough to make timely strategic decisions?
3. What are the implications of delegation, predelegation for contingencies, and unilateral lower-level assumption of authority?
4. What are the implications of normal decentralization of authority for control of events relevant to escalation? In other

* The unclassified literature on strategic C³I varies widely in accuracy and quality. For an apocalyptic and influential essay, see Ref. 7; for an interesting (but not always accurate) survey, see Ref. 8; for guarded discussion by knowledgeable experts, see Ref. 9. See also Refs. 10 and 14.

words, will the separate commands (CINCs and Soviet TVD commanders) be taking what they regard as standard measures that might on the one hand raise the likelihood of escalation or on the other hand fail adequately to anticipate the requirements of nuclear conflict or make nuclear strikes more lucrative? (Examples: ASW operations, threatening SSBNs or dispersal of nuclear weapons in conventional conflict, or--on the other end--massing of forces to achieve improved force ratios in conventional operations.)

5. What operational constraints narrow the NCA's employment options (e.g., concerns for SSBN vulnerability, weapon range, retargeting inflexibility, option purity, limitations in assessment capability)? What doctrinal constraints similarly narrow his options (e.g., failure to train crews for massive retargeting)?
6. What Soviet actions should we anticipate early in conflict by virtue of Soviet doctrine's emphasis on preparing for the nuclear phase? How should we prepare to observe, understand, protect against, and react to such measures?
7. What if we lose some or all communications (one-way or two-way) to the ICBMs, SSBNs, bombers, and/or CINCs...? Could we lose the capability for assured retaliation, limited responses,...? Is the EMP threat real and potentially devastating? What if we lose communications to the Soviets? Will that preclude termination short of unrestrained general nuclear war?
8. What if we lose early warning satellites from antisatellite attacks, sabotage on the ground, system failures, or unknown reasons?
9. Assuming the potential NCA desire to make limited responses to nuclear attacks, what capabilities at what level will be necessary to make appropriate limited responses possible--not only at the outset of conflict, but as a function of time thereafter?
10. What if the results of initial conflict are sharply different from those anticipated--because of system failures, surprise tactics, or whatever? What capabilities are needed to permit at least modest replanning?
11. What if communications are adequate initially but rapidly degrading? What are the effects on crisis stability and future ability to prosecute the conflict?
12. To what extent is a counterforce war made infeasible by fragility in the command and control system? What are the implications of one or both sides having command and control inadequate to prosecute an extended conflict?*
13. To what extent does the nature of actual war plans circumscribe the feasibility of limiting the scope of nuclear war, once it begins? Are they so inflexible as to preclude controlled responses or are they in fact adequately flexible given the limited number of plausible options, the relative predictability of certain aspects of strategic nuclear war, and the difficulty or cost of achieving enduring command and control to support more fine-tuned responses?

* Interestingly enough, it is sometimes argued that enduring C³I capability would be destabilizing, a view that would shock most defense professionals. See, for example, Ref. 10.

The list reinforces the conclusion that dressing strategic C³I means addressing an enormous range of issues involving everything from standard operating procedures at the operational level to the implications of satellite vulnerability. Upon a moment's reflection it is also evident that the answers to the questions posed depend upon such diverse variables as: (1) nature of the superpower leaderships; (2) succession and devolution of command; (3) locations and levels of conflict; (4) for history (initiating scenarios); (5) status of forces by type and theater; (6) range, quality, and availability of preplanned options; (7) the enemy's overall strategy; (8) technical issues such as the survivability of many systems (or functions); and (9) constitution capability for each component of C³IW. We shall begin to address them in Sec. IV.

IV. Conceptual Architecture for an Approach

Basic Considerations Governing Approach

Section III demonstrates that handling C³I within the structure of an ordinary modeling approach is simply not feasible--too many of the issues are less quantitative and naturally analytic than operational or behavioral, and it would be fruitless to try reducing a problem at hand to a very small number of simple variables. By contrast, the emerging capabilities of automated war gaming will be an ideal vehicle for exploring many of the issues systematically. Indeed, recognition that gaming was probably essential in creating effects such as command and control underlay much of the initial government interest in the RSAC project.[11] The RSAC project is now far enough along so that making this idea a reality is a high-priority item.

The most important premise governing our approach to strategic command and control is that such issues should be reflected in the very fabric of the RSAC system--to view C³I as merely one more effect for which a program "module" needs to be developed would be to misunderstand utterly the nature of the problem. Indeed, it would be closer to the mark were we to say that the various and sundry RSAC models should be embedded within the fabric of a command-and-control construct than vice versa.

All this implies that we need a conceptual architecture for our approach rather than a mere grab bag of physical models and artificial intelligence techniques. We shall now sketch out what we see as design requirements for our effort and then provide the outline of our intended approach.

Design Requirements. Upon reviewing the state of the art in strategic analysis, the major strategic issues of the day, as we see them, the nature of the RSAC charter from DoD,[1] and the conclusions of some past DoD work,[12] we have developed the following principles as guidance:

1. As mentioned above, the command-and-control issue should be reflected throughout the fabric of the simulation and not merely in some "module."
2. By virtue of our strategy-level focus, the approach should be top-down rather than bottom up. This implies we should be focusing on C³I functions rather than individual systems; moreover, it means that the key game variables and displays should be aggregated and in a form natural to strategic-level discussion.
3. The character of the system must account for the existence of multiple levels, and locations within levels, of command and

control authority. Moreover, it should reflect hierarchical phenomena.*

4. Even our early efforts to reflect command and control should be useful and realistic. It is better to reflect some of the real command and control issues early than to treat command and control comprehensively for a "toy problem" of no direct value.
5. However, the approach should be evolutionary and should allow linkup to some of the detailed work being conducted within the defense community on such matters as communications connectivity.
6. Although an evolutionary approach is appropriate, it should be broad-based from the outset--touching insightfully upon both U.S. and Soviet command, control, communications, intelligence, and warning rather than dwelling exclusively on, for example, U.S. capability to communicate an Emergency Action Message.
7. The approach must permit the analyst (or game director when human teams are involved) to introduce phenomena representative of the fog of war--phenomena such as may arise from causes as diverse as unconventional warfare, flukes of nature, or catastrophic weapon-system failure.**

Elements of Architecture. To move from requirements to an architecture we must first think about what would constitute an architecture. How would we know if we had one? Remember here that we are not dealing merely with the design of a communications system. Rather, we are dealing with an approach to the design and application of an interactive war-game-based simulation. Upon reflection, and upon looking back at Table 1, which itemized the elements of the RSAC automated gaming system, it seems we must provide four different items:

- o A suitable structure for the simulation (one that will provide the appropriate perspective on a multifaceted problem).
- o Variables (and corresponding data structures) suitable for reflecting command-and-control factors simply and transparently in rules and algorithms.

* The hierarchical principle of complex systems is what underlies the frequently mentioned analogy between living systems and command and control. Each level of a hierarchy has a recognizable separate existence and a set of internal processes. It communicates up and down the hierarchy, but the communications--however important--represent only a small fraction of the activities and are of little concern to most components of that level. Moreover, communications can be delayed, imperfect, inappropriate, misunderstood, etc. Manifestations of hierarchical effects are familiar to students of organizational theory, large business and government operations, biologists, and certain philosophers such as Arthur Koestler who see hierarchical principles as having broad applicability. We should note that it is one thing to build nested multilevel models, which are by no means unusual, and quite another to reflect hierarchical principles adequately.

** We wish to achieve here some important features of manual gaming in which the Control Team can force teams to focus on events that are systematically left out of typical model-generated scenarios. For example, a Control Team can decree to the Blue Team that its early-warning capability has vanished, and thereby force the Blue Team to think out what it would do in such a situation. The Control Team does not have to explain in detail how the catastrophe occurred (although good game practice would entail a plausible explanation). See Ref. 3.

- o Algorithms and rules for manipulating variables and data structures to model desired effects.
- o A management roadmap for assuring that C³I issues are introduced consistently throughout the complex RSAC system.

We shall now discuss our intended approach and touch upon each of these items in turn.

Outline of a Conceptual Approach

Simulation Structure. We shall not discuss the issue of simulation structure in much detail here because it should be reasonably clear from Fig. 1 and the attendant discussion that we anticipated most of the structural issues from early-on in the program. In particular, the RSAC's basic system design is well-suited to treatment of hierarchical and otherwise multilevel effects; moreover, it is flexible, modular, and designed with the expectation of evolutionary development from simple rule-based models to more sophisticated models using results of detailed work on such problems. Although it will be some time before we make extensive use of the capability, the system design also permits us to maintain separate data bases for Red, Blue, and "Real World." That is, our data structures permit us to have Red, Blue, and Force Agent to see different data bases on, for example, the status of Blue's forces. Handling such effects is painful for the analyst but will be essential for looking into matters such as deception and the fog of war.

Variables and Data Structures. Discussions of rule-based models sometimes seem to suggest that rule writing is easy and that structure is unnecessary--all that is necessary is to find an "expert." In fact, however, there are many instances in which experts able to provide a complete and incisive set of rules simply do not exist. In that case, which generally applies to command-and-control issues, it is necessary for analysts to do a great deal of background work to help experts think clearly and cover all the bases.[4] This implies anticipating (to the extent possible) what the experts will eventually find to be the most natural way to express rules simply and understandably from a strategic-level perspective. That is, we must anticipate the appropriate variables, related data structures, and logic flows.

Analytically, a major problem here is that the "natural variables" for those building the individual pieces of the RSAC simulation are often not the natural variables in which to express particular rules. For example, the RSAC has a world data base with information on worldwide forces, national orientations, etc. The data are collected initially in forms driven by the models that track locations and status of individual forces and the like. However, this form is too disaggregated and disorganized relative to what we need either to write simple decision rules involving command and control, or to write simple rules or models describing command-and-control effects in the execution of options.

With these considerations in mind, we are currently working out on a classified basis the details of an approach outlined in Table 2. Some of the basic notions here are as follows:

- o It is useful to construct three time-dependent state vectors in addition to O, which characterizes the "world data base" as it is evolving in RSAC work. The three vectors, C (with subvectors for the C³I and W components of C³I), and S, and N, pertain to the states of C³I, functional support for the NCA, and NCA capabilities, respectively.

Table 2

ORGANIZING DATA FOR SIMPLIFIED RULE-WRITING SENSITIVE TO COMMAND AND CONTROL

O + C + S + N			
State of Basic RSAC System + Observables	State of Command and Control	State of NCA Support (Functional Capabilities)	State of NCA Capability
O(t) = (unstructured data on: status of forces and nations' war plans (scripts) being implemented, attrition rates, rates of movement,...)			
C(t) = (C ₁ , C ₂ , C ₃ , I, W)			
C ₁ (t) = (Nature of NCA, Extent of Delegation; Extent of Contingency Prerelegation; Nature of NCA Staff; Degree of Information Saturation)			
C ₂ (t) = (Lower-Level Capability to Respond (to higher-level orders), Lower-Level Willingness to Respond (to higher-level commands))			
C ₃ (t) = (Communications (by geographic region) to: ICBMs, SSBNs; Bombers; SLCM launch platforms; Satellites; ASAT systems; Other Strategic Defensive Forces (SAMs, ABMs, interceptors); Nonstrategic CINCs)			
I(t) = (Intelligence on: Nature of enemy NCA; Nature of enemy NCA support; Enemy ICBMs, SSBNs, SLCM launch platforms, bombers, satellites, ASAT system, and Other Strategic Defensive Forces (SAMs, ABMs, interceptors; and enemy forces in theaters)			
W(t) = (Warning of attack by: Ballistic Missiles; Air Breathers)			
S(t) = (Ability of the NCA to Obtain Finished Assessments of: status of his forces, force operations, and alliances; status of the enemy NCA, NCA support, forces, and alliances. Ability of the NCA support staff to Use, Develop, and Evaluate Options, both before and after execution. Ability of the NCA's forces to Execute Options.)			
N(t) = (NCA ability to: Assess Option Feasibility; Modify or Originate Options; Compare and Choose Among Options; and Communicate the Chosen Option.)			

Transformations between these vectors should be thought of as transforming raw data on system observables into forms more convenient for rule writing.

- o Each of the state-vector components and their time trends should be definable, for our purposes, in highly qualitative terms. For example, we may characterize the quality of the first component of C₃, communications to ICBMs in a given geographic region, as: poor, moderate, or good.
- o We would expect to write nearly all rules involving the command-and-control influence on option selection in terms of the vector N (and C₁, which determines the major agent's character); other rules, however, (e.g., Force Agent rules on option execution) may depend on S, C, or--in rare instances--information found only in the world data base.
- o In defining the state vectors, their components, and the values of their components, we must be cautious to maintain resolution distinguishing among the following, even in early work:

- Theaters: intercontinental, space, others (Europe, SWA,...), and simultaneous multitheater operations
- Strategic Forces: ICBMs, SLBMs, bombers, SLCMs, Space Forces, and ASAT Forces
- Time: crisis; extreme crisis and possible theater war; period of U.S. first strike; period of

- Soviet first strike; period of immediate U.S. response, if any; initial aftermath; and extended aftermath (see Fig. 2, but note also the possibility of more complex stop-and-start wars)
- Option Class: e.g., limited versus massive counterforce options with modest or major coordination problems (including theater missions for strategic forces) limited and massive countervalue options, and mixed options, in each case executed as a first-strike, Launch Under Attack, Prompt second-strike, Delayed second-strike, or Follow-on strike.
 - Employment Concept: distinctions among options calling for the same results to be achieved with different missions for the individual force types (e.g., striking the same targets with bombers as opposed to ICBMs).
 - Class of Effect: effects on ability to choose, quality of choice, and speed of choice.

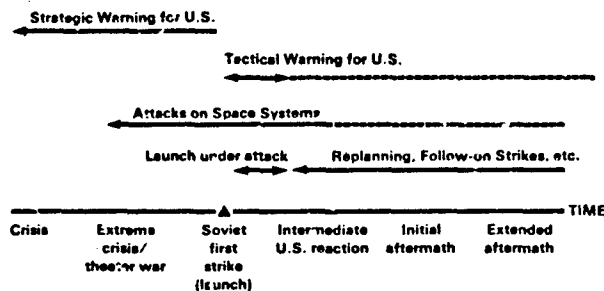


Fig. 2--Illustrative Time Periods and Events

- o The NCA's ability to perform the functions listed in N(t) will depend strongly on the types of options under consideration. Although options can theoretically be indexed by all components of the preceding bullet, we thus far believe that Time and Option Class are the most important characteristics for option indexing. Thus, the NCA's ability to assess option feasibility, for example, should be understood to be with respect to each Option Class component within each Time component. For each such combination we are now in the process of defining High, Medium, and Low labels for the components of N(t).

Even further distinctions will clearly be necessary when we begin to delve more into theater-level issues, but the above list is already intimidating.

Although the many distinctions itemized above may seem to imply that aggregation can't work, in fact there appear to be many possible simplifications. For example, in characterizing the capability of the NCA to develop new options during the "Period of Immediate Response," we might use Low, Low-Medium, High-Medium, and High as the principal descriptors with these definitions:

Low: No capability except for execution of preplanned options in class X (X to be defined in terms consistent with the war plans/scripts available). No

- retargeting. No capability to change theater war plans.
- Low-Medium: Preplanned options of classes X and Y with retargeting of force elements A and B feasible (where A and B are the force elements for which retargeting is most plausible). No capability to change theater war plans from the NCA level.
- High-Medium: As above, except full retargeting across force types within preplanned options.
- High: Full nominal capabilities as of some future date, including in particular the ability to retarget against newly acquired targets in ad hoc options.

These definitions allow a few simple "values" to cover a number of issues. We would also need separately to characterize the components in terms of timeliness of decision (e.g., normal or slow).

Obviously, the approach involves an article of faith to the effect that the strategic command-and-control problem can be reduced to describing capabilities and phenomena in a large but highly finite number of crudely defined discrete states, preferably states that can be summarized briefly in intuitive terms. Considering that most strategic nuclear analysis implicitly assumes the state of perfect command and control (except for zero strategic warning), we need hardly apologize for an approach that will distinguish among tens (or perhaps hundreds) of states. How much disaggregation will be necessary remains to be seen.

Roadmap to Integration. Assuming that the structure we have outlined provides an appropriate view of the problem, and that its states are defined for rule writing and model building by using the natural variables of the command-and-control problem, the next challenge is to manage the implementation. Unfortunately, this is inherently difficult because, as

repeatedly stated, C³I permeates everything and must therefore affect the work of numerous people working on different parts of the RSAC project. There are at least three aspects to managing the work in such a case: (1) rule writers and model builders must have checklists of items to consider, thereby reducing the likelihood that Red will write rules sensitive to some command and control issue that Blue will ignore, except in those cases where underlying strategic asymmetries dictate valid differential sensitivities; (2) there must be a mechanism of integration in which the various contributors systematically read each other's material, compare notes, and look for incompatibility; and (3) there must be formal "walkthroughs" of the overall simulation on every command and control issue expected to be important.

We cannot discuss these matters in much more detail here, but we can point out a few items of interest. For example, upon reflection we find it useful to distinguish clearly between command-and-control effects on decisions, and command-and-control effects on force operations. Figure 3 makes this distinction and points out that all of the RSAC agents are affected. Note that:

- o The state of strategic command and control (and the projected state!) must affect Red and Blue decision rules by: (a) affecting Red or Blue character, warfighting ability, and efficiency; (b) limiting Red or Blue options; and (c) shading the perceived attractiveness of alternative available options. Similarly, Scenario Agent's decision rules must be

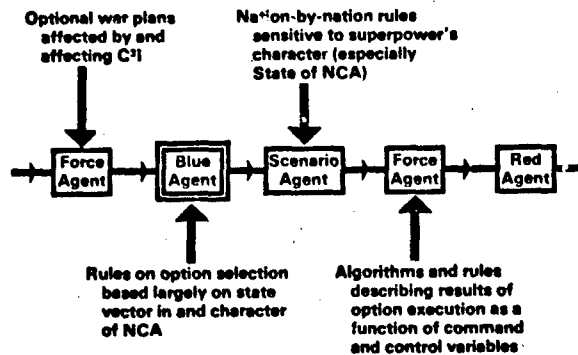


Fig. 3--Command and Control Effects on Blue Move and Results Thereof

sensitive to the nature of the superpowers' national command authorities and the overall effectiveness of those superpowers.

- o The Force Agent must reflect command-and-control effects on force operations (and of projections thereof) by means of: (a) delays and related mismatches between decisions and current world states; (b) errors such as those caused by poor intelligence or communications; (c) degraded capabilities such as loss of certain types of warning or intelligence; and (d) coordination problems.
- o The analytic war plans/scripts must reflect to some degree the partially independent operations of individual theater commanders and the potentially parochial decision rules governing those operations.* They should reflect doctrinal behavior at the operational level except where there are good reasons to assume otherwise.**
- o The analyst using the RSAC system must be able to insert "fog-of-war" effects and other related phenomena easily, something that has implications for System Monitor in particular, but also for the other Agents (i.e., there must be variables created to serve as surrogates for the effects in question; the variables must be represented in all of the separate agents).

All of this is rather abstract, so it is useful to provide at least a partial image of what is involved in implementing the concept. Thus, let us discuss what might be involved in reflecting just one particular issue, Red's assessment of Blue's LUA capability. Such an assessment might be important in Red's detailed attack planning if the implications of a U.S. LUA were major. Figure 4 suggests a somewhat oversimplified logic.

To implement this logic in the RSAC system one would have to do the following:

* Ultimately, we hope to reflect independent operations by commanders at levels lower than the theater. In the relative near term, however, we will omit such considerations.

** This has management implications because it suggests that we should invest in having separate teams develop the war plans for the individual theaters rather than building the plans from a purely top-down perspective that would tend to make the analytic plans used in the computer model come out far more coordinated and mutually reinforcing than is realistic. Unfortunately, developing such separate plans is manpower and expert intensive, especially because of the effort required to train teams of analysts.

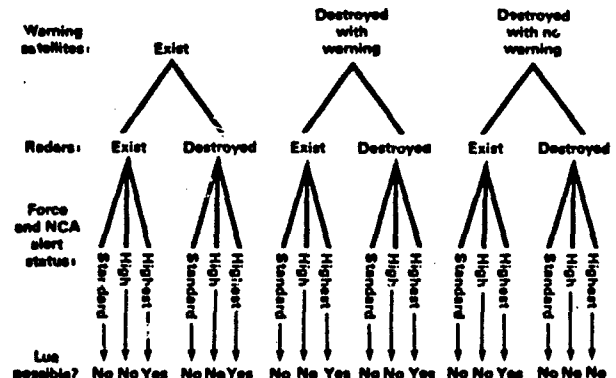


Fig. 4--A Simplified Logic Tree for Red's Assessment of Blue's Ability to Launch Under Attack

1. Create Red war plan components that would, if successful, destroy Blue's warning satellites and radars.
2. Build Force Agent models to estimate the effectiveness of such Red options under a variety of circumstances (e.g., the status of Red's antisatellite systems and space-tracking network, the number and vulnerability of U.S. satellites, the locations of Red's SSBNs capable of attacking U.S. warning radars).
3. Build Red decision rules sensitive to Red's assessment of Blue's LUA capability (e.g., ruler affecting Red's willingness to launch a first strike or rules affecting the size and nature of a first strike; also, rules relating Red's desire to prevent U.S. LUA to other Red actions that would provide the United States with strategic warning).
4. Build Red rules assessing Blue's LUA capability (as in Fig. 4) and relating the items in the figure to the war plan components mentioned in (1) and Blue's probable response to other Red war plan components as mentioned in (5).
5. Build Blue rules sensitive to indication that warning satellites are under attack (e.g., go on highest alert) and rules establishing whether Blue would actually try to launch under attack under some circumstances (a function of policy and capabilities, which might be quite different from those assumed by the Soviet Union, whose strategic doctrine has stressed LUA for years).

Although this is only a narrative sketch, it is sufficient to demonstrate once again that incorporating command-and-control effects is an inherently complex business demanding that attention be paid to details of scenario, strategy, the two-sided nature of the game, etc.

Where, then, do we stand at this point in our development program? Is this all conceptual, or are we actually implementing the ideas? At the moment, we are within a few months of automating the most recent version of the basic RSAC system, having conducted semiautomated experiments last summer.[13] Once the basic system is operational, we plan to incorporate selected command-and-control effects on a simplified basis using heuristic rules tied to grossly defined world states (e.g., have the Soviets already attacked warning satellites?). We then expect to implement a more ambitious but still first-generation of the overall architecture, probably in November or December, 1983. Finally, we expect to build more sophistication into the system over the period of several years--

including explicit tie-ins to the results of detailed models such as those used to estimate connectivity to bombers as a function of weapon lay-down and scenario. We plan to use structured human gaming as a source of insight and rules. Our expectation is that applications will be possible early next year, well before we have much sophistication--primarily because a major contribution of the effort will be a war game framework requiring consistency from move to move and requiring the human or automated players to take first-order C³I effects into account when developing their overall strategies.

References

1. Davis, Paul K., and James A. Winnefeld, The Rand Strategy Assessment Center: An Overview and Interim Conclusions About Potential Utility and Development Options, The Rand Corporation, R-2945-DNA, March 1983.
2. Steeb, Randall, and James Gillogly, Design for an Advanced Red Agent for the Rand Strategy Assessment Center, The Rand Corporation, R-2977-DNA, May 1983.
3. Davis, Paul K., Concepts for Improving the Military Content of Automated War Games, The Rand Corporation, P-6830, November 1982.
4. Davis, Paul K., "Concepts for Managing Complexity in Realistic Applications of Artificial Intelligence to Problems of Strategy and Operations," The Rand Corporation, forthcoming.
5. Glaser, Charles, and Paul K. Davis, Treatment of Escalation in the Rand Strategy Assessment Center, The Rand Corporation, N-1969-DNA, March 1983.
6. Stan, Peter, and Paul K. Davis, "Escalation in Automated War Games," The Rand Corporation, forthcoming.
7. Steinbruner, John D., "Nuclear Decapitation," Foreign Policy 45, Winter 1981-82.
8. Ball, Desmond, "Can Nuclear War Be Controlled?" Adephi Papers 169, International Institute for Strategic Studies, 1981.
9. MITRE Corporation, National Security Issues 1981 Symposium, MITRE Document M82-30, 1982 (reporting on a symposium held by the Air Force Electronic Systems Division and MITRE during October 13-14, 1981 on the subject of "Strategic Nuclear Policy, Weapons and the C³I Connection").
10. Tucker, Johathan, "Strategic Command-and-Control Vulnerabilities: Dangers and Remedies," Orbis 26, No. 4, Winter 1982.
11. Marshall, Andrew, "A Program to Improve Analytic Methods Related to Strategic Forces," Policy Sciences Vol 15, No. 1, November 1982. See also Appendix 1 of Ref. 1.
12. Harris, Donald, Andrew Marshall, Eugene Durbin, and Fred Giessler, Final Report of the OSD Research on Military Command Control, 1979, Office of the Director, Net Assessment, Department of Defense, 1979.
13. Winnefeld, James A., Illustrative Experiments with an Interim Version of Rand's Strategy Assessment Center, The Rand Corporation, N-1917-DNA, November 1982.
14. Van Creveld, Martin, Command, Unpublished report for the Director of Net Assessment, Office of the Secretary of Defense (n.d.).

AD P002876

A COMPARISON OF TWO AIR DEFENSE COMMAND AND CONTROL MODELS

Dr. John R. Dowdle

Dr. Leslie C. Kramer

Mr. Robert F. Gendron

Mr. Mark P. Merriman

ALPHATECH, INC.
111 Middlesex Turnpike
Burlington, MA 01803
(617) 273-3388

SECTION 1 INTRODUCTION

In this paper we describe and compare two models of air defense systems which focus on the command and control (C²) aspects of the air defense mission. The models are both computer-based representations of how, through the execution of C² functions, a defensive force detects and destroys an offensive force of airborne penetrators. However the two models are completely different in the manner in which this engagement is represented. The first model, called QUEB for queuing based, consists of a large set of simultaneous equations derived using queuing theory which solve for the dynamic steady-state conditions existing throughout the system. The second model, called TADZ for transient air defense zone, is a so-called event-stepped simulation of the air defense engagement.

We first describe an underlying skeletal or conceptual approach to modeling command and control; this skeletal approach is reflected in both QUEB and TADZ. We then describe each model in turn. Finally, comparisons are drawn between the two approaches to modeling the same problem.

The contribution of our work is two-fold. First, the underlying approach to modeling C² processes we believe to be generic and applicable to other command and control missions. Second, the illustration of the generic approach in terms of two specific models which have actually been implemented and run gives insight into the spectrum of conclusions one can draw (using different sorts of models) about how command and control contributes to the achievement of a military mission.

SECTION 2

C² SKELETAL MODEL AND MODEL ELEMENTS

2.1 Basic Approach

Command and control is, in some sense, the "glue" that holds together a complex collection of military hardware and people trying to use it for a purpose. To achieve a military mission, the typical problem is to apply scarce resources in a timely manner to meet a threat (defense) or achieve an

objective (offense). For example, in air defense the problem is to detect threats and assign weapons to shoot them down prior to their carrying out their assigned attack. The scarce resources in the air defense environment include the obvious ones such as fighter-interceptors (FIs) or surface-to-air missiles (SAMS). Other resources may be more subtle, such as "attention slices" that radar operators must devote to penetrators to recognize their presence, identify them, and generate reports for subsequent processing by the system.

Systems characterized by competing demands upon scarce resources, each of which is occupied for a finite time while it is being used, are modeled mathematically within the discipline of queuing theory [1]. We therefore argue that a queuing-based formulation of command and control is an appropriate one. We shall illustrate such a formulation below.

A second aspect of command and control that must be captured by any model is that C² systems are feedback (or "closed-loop," or "cybernetic") systems. This is illustrated in Fig. 2-1. The generic C² system consists of sensors which gather information, decision makers who act on this information, and weapons systems which are directed on the basis of commands issued by the decision makers. All this is tied together by a communications network of some sort which passes the information around. (In our view, some analysts mistakenly consider the communications network alone to be the "C² system.") The system is closed loop because the sensors see what happens! The feedback nature of the system is critical to its evaluation because the load on the system is governed by the need to re-engage targets which are not successfully destroyed the first time they are processed. It is obvious, then, that models of C² systems must include models of the sensors, decision makers, and weapons associated with a mission, as well as a model of how they are tied together.

In our view, the appropriate measures of effectiveness (MOEs) for characterizing C² relate to time, leakage, and resource utilization. How long does it take to prosecute the average target? How many targets escape? How often are there no FIs available to engage known threats, or how often are all radar operators busy, causing

new targets to go undetected? Queuing formulations deal precisely with these questions, as we shall show below.

2.2 C² Element Models

Given the structure described above, the analyst must determine the degree of detail needed in the individual element models inserted into the structure. Our objective in the work reported here was to include enough detail to address the questions for which the model was intended, yet as little detail as possible about sensors, decision makers (DMs), and weapons. An explicit objective of the model design was to maximize execution speed, and one means of achieving this was to include only those model details actually needed.

SENSORS

The sensors in our models are represented by single-look probabilities of detection. This parameter may be influenced by geography and relative sensor/target location, e.g., in terms of terrain masking. It may also be influenced in some cases by offensive action such as jamming. The analyst using the model is able to implicitly describe detection probability in terms of physical quantities such as radar power and target cross section, which are then passed through the radar equation to determine detection probability. However all of these physical parameters influence the C² computations through the one probability number.

The only sensors we have modeled to date are radars of various types, both ground based and airborne.

WEAPONS

As with the sensors, we wished to minimize detail in our weapons models. Thus, we chose to represent weapons in terms of one-on-one probability of kill. Other parameters associated with weapons in more detailed analyses of weapons effects, such as availability, for example, are "rolled into" the single engagement kill probability. Multiple engagements of a target by one or more defensive assets are characterized by the appropriate chaining of one-on-one kill probabilities.

The number of weapons of a given type and their geographic location are explicitly represented in the models; it is possible for all weapons in a region to be "busy" prosecuting penetrators with the result that subsequent penetrators pass through unengaged.

DECISION MAKERS

We draw a distinction between decision making tasks which are "easy" and those that are "difficult." Easy tasks are those whose execution can be represented merely in terms of a delay or in terms of an algorithm in combination with a delay, the algorithm specifying which of several outcomes results

under given conditions. Difficult tasks are those requiring "intelligence" or "insight" to carry out.

In our models to date, we have represented only easy tasks, i.e., DMs are represented as delays and algorithms. We have not yet considered model applications requiring modeling of cognitive processing or "intelligent" behavior.

INFORMATION NETWORK

The information network we have modeled to date is considered to be divided into a surveillance side and an airspace control side, as shown in Fig 2-2. This division is for conceptual clarity only; in reality, surveillance side and control side installations in an actual air defense system may be physically colocated. Both surveillance ("S-side") and airspace control ("C-side") processing is represented as a three-level hierarchy. SOs are the basic target detection and report generation units. These report to SIs, which are fusion centers, themselves reporting to a single S2 which carries out global data fusion for the air defense system. (Some SOs may report directly to S2 if desired.) On the control side, a single C2 divides target assignments among several CIs, which further subdivide them among individual COs which represent SAM or FI units. There may or may not be cross-links from SOs to COs or from SIs to CIs; the presence or absence of such links results in effects which we wish to study with our models, as we shall show below.

In this section, we have described the generic features of our approach to C² modeling. We now describe how these are reflected in the two specific models constructed to date.

SECTION 3

QUEUING-BASED MODEL (QUEB)

In this section, we will describe the features of the QUEB model, its formulation, and the computations carried out when it is executed. This model is documented more completely in references [23]-[4].

3.1 Model Features

The air defense engagement modeled by QUEB consists of the interaction of a homogeneous threat with a mixed defense. That is, the threat is assumed to consist of a single type of penetrator, say B-52s or cruise missiles, while the defense is assumed to include both SAMs and FIs (assuming the model is used to represent a US attack on the Soviet Union.) The threat is assumed to traverse a number of discrete penetration paths which are fixed a priori but which are unknown to the defense. Similarly, the ground based defense deployment is fixed throughout the engagement. (Airborne defense elements move about based on their assignments to targets.) The penetrator routes may be chosen by the model's user to avoid defense assets if desired.

The information flow represented is that illustrated in Fig. 2-2, however the cross-links from SOs to COs and from SIs to CIs are absent. Thus the information flow is assumed centralized in that all target reports pass through S2/C2. Message generation on the surveillance side is probabilistic, representing non-perfect single-look detection capability, however the detection model is a simple "cookie-cutter" approach. That is, each SO is assumed to "cover" a particular geographic region associated with which is a single probability of detection per look number. Radars can achieve a number of looks governed by the penetrator speed, the radar sweep rate, and the penetrator route relative to the radar position.

The control side represents several types of threat prosecution, depending on whether the particular unit carrying out an engagement is a SAM, a single FI, or a group of FIs (a "flight"). As target assignments percolate downward from C2 to CIs to COs, the messages are handled differently depending on the type of defensive asset involved. For example, FIs are assumed to have on-board radars (of characteristics specified by the model user) which must acquire the target before the FI can engage it. SAMs, on the other hand, must be directed by ground-based fire control radars. Flights of FIs can be represented as jointly prosecuting targets.

In the QUEB model, the DMs are modeled simply as delays. This is adequate since alternative actions on the part of the DMs are not modeled: it is clear from the information network model topology where all information goes. The DMs have only to recognize the presence of targets (DMs at SOs), carry out fusion (DMs at SIs and S2), and make increasingly detailed target assignments (DMs at C2, CIs, and COs). While these are perhaps stressing assignments in real life and require some "thought," it is adequate to represent them merely as delays in QUEB because the outcome IN TERMS OF WHO DOES WHAT NEXT is highly constrained -- although target assignments on the control side are reactive to engagement conditions, the fundamental information flow pattern SO-S1-S2-C2-C1-CO is fixed. (Looking ahead, this is not the case in TADZ.)

The C2 system structure is fixed in any given QUEB run. That is, neither the topology of the information network nor the characteristics of the system elements changes during an execution of the model.

3.2 QUEB Formulation

QUEB is a set of many simultaneous equations, derived from queuing theory [1], which are iteratively solved as the model is executed. Queuing theory deals generically with the set-up depicted in Fig 3-1. Here "customers" arrive into a queuing module, which attempts to place "servers" at their disposal to carry out some task. A customer which arrives when all servers (of which there may be one or more) are busy waits in

a "queuing facility" until either a server becomes free to attend to it or the customer gives up and leaves the module (in which case it is said to "renege"). These terms and concepts become intuitively clear if the reader considers entering a fast-food restaurant seeking a meal.

Queuing theory permits an analyst to compute several statistics about the set-up depicted in Fig. 3-1. Given statistics characterizing the arrival of customers, statistics characterizing the time it takes to process each one, and the number of servers, the theory provides equations describing the statistics of the waiting time that passes until a server is available to process a customer, the statistics of reneging (i.e., how often customers leave the module without being served), and the statistics of server utilization (i.e., how often everyone is busy.) Examples of such equations are not given here due to space limitations; the interested reader should see [1]-[3].

Within the QUEB model, the elements of the C2 system are identified with the queuing concept of a server, and the penetrators and messages about them are identified as customers. Thus, for example, any SO is considered to be a server which is trying to service a penetrator by detecting it and generating a message which is a target report. The operators at the SIs are trying to service the incoming target reports by fusing them, i.e., by determining which reports from separate SO sites represent in fact the same physical target. The individual threat prosecution assets (FIs and SAMs) represented at the CO level are attempting to serve the penetrators by locking their local sensors onto them and destroying them. Penetrators (again represented as customers) are, from the viewpoint of their mission, only successful if the defense fails to complete their service before they leave the airspace of the air defense system. Thus, the penetrators that leak through the defense are, in queuing terminology, the reneges.

The overall structure of the QUEB model is shown in Fig. 3-2. There is a set of equations representing the surveillance side of the system which calculate the statistics of message generation. There are equations representing the resource allocation modules which compute the statistics of order generation. Finally, equations representing the threat prosecution modules compute the statistics of as yet undestroyed targets. These last feed back to the surveillance modules as shown in the figure. Equations associated with all of the modules represent the overall measures of performance and effectiveness associated with the system.

3.3 QUEB Computations

As stated above, QUEB consists of a large set of simultaneous nonlinear equations from queuing theory. These are solved iteratively to yield the dynamic steady state statistics throughout the system. These statistics can be interpreted

as the system leakage, the percentage of asset utilization, etc.

SECTION 4

TRANSIENT AIR DEFENSE ZONE MODEL (TADZ)

In this section we will describe the TADZ model in a manner parallel to the discussion of QUEB above: we will cover the model's features, its formulation, and the computations necessary to execute it. TADZ is documented in detail in references [5] and [6].

4.1 Model Features

The primary motivation behind the design of TADZ, which was the second air defense model we developed, was to provide a capability to explicitly determine the transient behavior of the air defense C² system. In addition to this desire, which we satisfied by going to an event-stepped simulation as we discuss below in Section 4.2, there were several mission features we wanted TADZ to deal with which were absent in QUEB. (QUEB could be extended to include one or more of these features if desired; see Section 5.)

The first feature desired was an ability to deal with a mixed threat, i.e., one consisting of several types of penetrators (bombers, cruise missiles, etc.) arriving simultaneously. This feature would allow study of the impact of different engagement priority doctrines as executed by the C² system. Coupled with this was a desire to model a flexible, adaptive C² structure which would permit information routing to be responsive to the progress of the engagement. In contrast to QUEB, which has a fixed, centralized C² structure, TADZ allows the structure to vary during a run both because the defense may choose to change the routing of messages and because the offense can destroy elements of the C² system. The C² system adaptivity is manifested by its ability to react to the destruction of defense elements and continue the engagement using a changed C² structure.

Certain additional mission elements were included in TADZ beyond those present in QUEB. These include ECM applied by the offense, defense suppression (as just mentioned), terrain masking of the defense radars and exploitation of terrain masking by the offense, and resupply of exhausted defense elements such as FIs which have used up their fuel and SAM launchers having no more missiles.

The DMs in TADZ are represented both as delays and as algorithms. The algorithm representation is needed because in TADZ the outcome of their decisions IN TERMS OF WHO DOES WHAT NEXT is not as rigidly fixed as it is in QUEB; operators on the S-side of the system can decide where to route messages, either upward through the nominal centralized chain of command or across to the C-side directly. The C² adaptivity

captured in TADZ requires that the DMs decide where to send messages as a function of the engagement conditions as interpreted locally. There are different classes of messages represented in TADZ, for example classes representing targets of different priority, which are processed differently as governed by algorithms representing C² doctrine.

The information network in TADZ is that shown in Fig. 2-3 including the potential cross-links from SOs and COs and those from Sis to CIs. The DMs in the system determine which links to route information over as described in the preceding paragraph.

4.2 TADZ Formulation

Queuing theory does provide techniques for transient analysis through computing the time evolution of the probability distributions governing the statistics of the variables in the system. However these involve solving partial differential equations. In our judgement, the computational effort required to do this was sufficiently greater than that required for solving the QUEB simultaneous algebraic equations that we preferred to examine alternative approaches. Alternative approaches were also considered because the project supporting the work reported here is explicitly directed at investigating diverse techniques for analyzing C². A second model of a second type was requested by the sponsor.

TADZ is an event-stepped Monte Carlo simulation. The C² elements and their interconnection are represented as nodes in a network, each node being represented as a queuing module as shown in Fig. 4-1. Any given node is preceded by others (or by a target generation process) and followed by others (or by targets exiting from the controlled airspace). As indicated in Fig. 4-1, messages or targets arrive and are interpreted as customers in a queuing sense. Activities are carried out at the node and messages are sent to subsequent nodes as appropriate, i.e., the customers are served. In the course of service, events are defined, for example, target enters radar coverage, target is detected, SAM is assigned, target survives engagement, etc. Events signal the start or completion of various C² processes, and the time necessary to execute these processes is either supplied as an input to the simulation by the user or computed according to a model embedded in it. The simulation advances time as event after event occurs; this is the sense in which it is event-stepped. In other words, simulated time does not progress in uniform increments.

As the status of each node in the network (as represented generically in Fig. 4-1) is updated as a result of events, statistics are collected for later display. Examples of these statistics are the number of busy operators at each node, the current backlog in the many queues, the time taken to process messages (for those having stochastic times), etc.

As stated, the simulation is a Monte Carlo model. Stochastic effects represented in the model over which the Monte Carlo samples can be averaged include the arrival times of the penetrators, the outcomes of detection opportunities, the outcomes of one-on-one engagements, and the time necessary to carry out various decision processes as well as others. Deterministic runs can be obtained by "turning off the noises."

4.3 TADZ Implementation

Monte Carlo simulations of C² systems can be very time consuming to execute. We developed a design for TADZ which keeps as much calculation as possible outside the Monte Carlo loop so that it is only done once.

The simulation's components are shown in Fig. 4-2. Input and output processors communicate with the user. An event processor executes outside the Monte Carlo loop to set up the potential events that might occur during the engagement, i.e., these are tentatively scheduled. The concept is illustrated in Fig. 4-3, which depicts a typical penetration route passing through a region of surveillance coverage and two "lethality regions" within which penetrators are vulnerable to SAM attack. Based on the penetrator speed, the nominal times when penetrators would cross the boundaries of these regions are computed and tentatively scheduled. Whether these events "actually" occur remains to be seen.

After the preprocessing just described is completed, the cyclical calculations are then carried out by the network processor using physical subsystem models and information flow models to compute the outcomes of many random events. Some of these modify the precomputed event schedule so that it reflects the experience of a single Monte Carlo run as it executes. In this sense, some events that were prescheduled do not "actually" occur. Each pass through the Monte Carlo loop begins with the same tentative event schedule, however each pass through the loop typically completes with a unique tableau of events which actually took place.

SECTION 5

COMPARISON AND CONCLUSIONS

Two models of air defense command and control, both built using an underlying conceptual framework for viewing command and control problems, have been described in this paper. Both have been implemented, and the results obtained using them have been compared.

Specific numerical results obtained with these models cannot be given in an unclassified publication. However it is useful to contrast the characters of the two models. The QUEB model is an analytic model in the sense that its execution involves the solution of a set of simultaneous equations. It is implemented in about 24000 lines of

PASCAL code and runs in a couple of seconds on a VAX 11/750. We have used it to study the impact on C² effectiveness of various technology options. We have seen, for example, that several typical surveillance and threat prosecution upgrades do not have a significant effect on an overall measure of effectiveness such as "average time between penetrator entry and its first encounter with a defender." It is clear why this is so when the numbers are examined -- in short, geographic and geometric effects dominate this MOE under the conditions assumed (and not necessarily under all conditions). However the system upgrades studied do have a significant effect on the pattern of loading imposed on the C² assets. For example, better surveillance leads generally to earlier detection with a concomitant increase in load on elements at the "front" of the air defense system, i.e., those along the edge over which the penetrators enter.

The TADZ model is an event-stepped simulation implemented in about 75000 lines of FORTRAN. The higher-level simulation language SLAM has also been used in the implementation of TADZ. It takes about 15-30 minutes to execute a typical scenario using a nominal number of Monte Carlo iterations, e.g., 10. Thus TADZ is as much as 1000 times slower than QUEB. TADZ does provide much more detailed information, however. It shows the analyst the impact of flexible C² doctrine and how the C² system reacts to perturbations such as destruction of important C² centers during an engagement. When factors represented both in QUEB and TADZ are compared, the two models have been found to give similar results.

We have concluded from our studies that alternative and complementary models of C² can be useful for studying various questions about C² systems. We have found that valuable insight can be gained about C² MOEs using models having minimal detail in them about sensors, weapons, and decision makers; it is the interconnection and interaction of these elements that matters most for some studies. Generally, the obvious is true: more detail causes a slower running model. The analyst must determine if the detail to be included is really needed for his purposes, or whether he is succumbing to the temptation to put it in because "more must be better." We point out that detail does not in itself imply that a simulation model is required in place of an analytic one. As we said in Section 4, several of the features included in TADZ that are absent in QUEB could have been added to QUEB if desired.

Finally, we have convinced ourselves that the skeletal framework for modeling C² described in Section 2 is a generally useful one applicable to more than air defense C² analysis.

REFERENCES

- [1] Kleinrock, L., *Queueing Systems*, Wiley-Interscience, 1975.

21-a

21-a

[2] Lauer, G.S., M.G.Bello, J.R. Dowdle, and J.C. Kramer, "C3 Systems Dynamics Project: First Annual Report. Vol I: Conceptual Framework and Preliminary Methodology. Vol. II: Air Defense Problem Study (U)," TR-139, ALPHATECH Inc., Burlington, MA, Sept. 1982, SECRET.

[3] Merriman, M.P., A.L.Blitz, M.G.Bello, and J.R.Dowdle, "C3 System Dynamics Project Software Report (U)," TR-140, ALPHATECH, Inc., Burlington, MA, Jan. 1983, SECRET.

[4] Merriman, M.P., "C3 System Dynamics Project Software Report: QUEB II," (Two Volumes) TR-182, ALPHATECH, Inc., Burlington, MA, October 1983.

[5] Dowdle, J.R., R.F.Gendron, W.S.Powell, and G.F.Fitzmaurice, "Transient Air Defense Zone Model (U)," TR-188, ALPHATECH, Inc., Burlington, MA, Nov. 1983, SECRET.

[6] Merriman, M.P., R.F.Gendron, and G.F.Fitzmaurice, "C3 System Dynamics Project Software Report: TADZ," (Two Volumes) TR-186, ALPHATECH, Inc., Burlington, MA, Nov. 1983

This work was supported by U.S. Air Force Systems Command under contract F33657-81-C-2150.




Fig. 2-1
C² System Elements and Interrelations




Fig. 2-2
C² System Information Network




Fig. 3-1
Queuing Module Elements

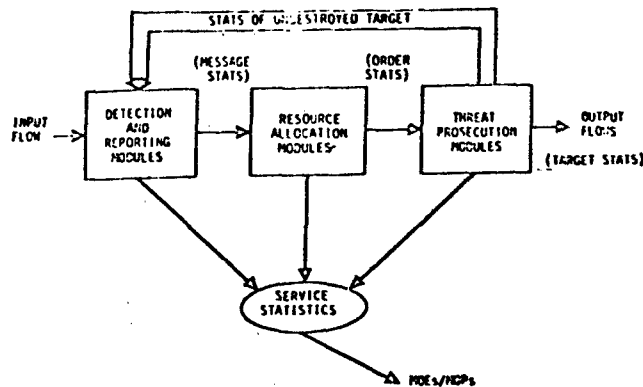


Fig. 3-2
QUEB Model Structure

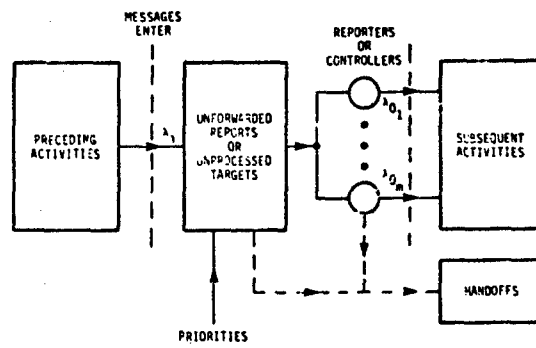


Fig. 4-1
Generic TADZ Node Model

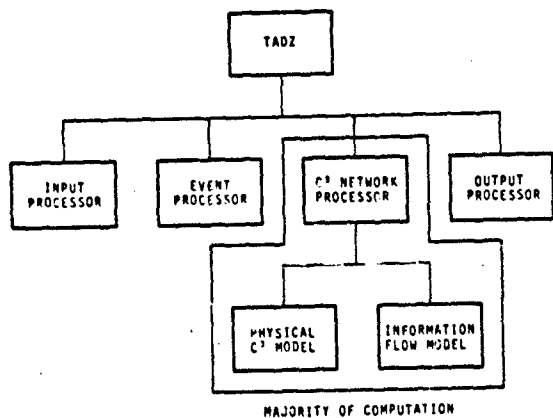


Fig. 4-2
TADZ Implementation




Fig. 4-3
Sample Event Scheduling

LANCHESTER'S EQUATIONS AND GAME THEORY

J. M. Hozencraft
P. H. Moose

Naval Postgraduate School
Electrical Engineering Department
Monterey, California 93943

SUMMARY

Generalization of Lanchester's Equations leads to a model of combat in which different types of Blue forces (say infantry, artillery, armor) suffer attrition due both to aimed and area fire from different types of Orange forces. The attrition coefficients depend on the types of forces and the conditions of combat. The problem considered here concerns optimization of the force distributions $x = (x_1, x_2, \dots, x_I)$ and $y = (y_1, y_2, \dots, y_J)$, subject to constraints on the aggregate forces ($\sum x_i$ and $\sum y_j$) and given the set of attrition coefficients. It is shown that choosing the objective function as the difference of the aggregate attrition rates leads to an optimization problem that is mathematically equivalent to a matrix game. It follows that the combat model has a saddle-point solution, with optimum force distribution vectors x^* and y^* which can be determined by a linear program.

The remainder of the paper investigates the neighborhood stability of the model at the optimum operating point (x^*, y^*) . We define and calculate two parameters, k_1 and k_2 , which partially characterize stability. A two x two example is presented along with sample unstable phase trajectories.

Command and Control decisions have great influence on the constraints governing the feasible set of different force compositions which can be brought to bear in a battle. It is hoped that better understanding of how these constraints affect the (optimum) battle outcome will provide insight into the quantitative impact of command and control information systems. Substantial additional work generalizing and interpreting the results to date will be necessary, however, before this becomes possible.

LANCHESTER'S EQUATIONS AND GAME THEORY

Background

Consider a battlefield and assume that Orange and Blue are opponents. Each has several different types of forces available - say infantry, artillery and armor - and the total forces available to each side are X and Y , respectively. Our problem is to determine an appropriate distribution of forces for each side, and to evaluate the cost of doing battle.

A more elaborate problem of the same genre would involve two or more battlefields, with individual constraints on the total strength of each type of force. We again wish to apportion the forces (across the battlefields) and evaluate the cost of battle.

Motivation to study problems of this sort stems from a belief that the development of a theory of Command and Control requires that we be able to quantify the effect of combat information processes on battle outcome. These information processes feed decisions which in due course constrain the set of feasible force distributions. A measure of the utility of information might therefore be taken as the change in the "optimum" cost of battle implied by a change in the constraints on what forces can be brought to bear.

Additional motivation arises out of the need for interactive wargaming as a vehicle for developing new command and control doctrine. Even in a "distributed" game, where most players can participate in situ, it will be necessary to use player surrogates and to aggregate combat results from one level of resolution to another. Hopefully, the models studied here will provide insight both into how to aggregate meaningfully, and how to design plausible surrogates.

Problem Formulation and Results

In order to model the combat processes we adopt the following generalized version of Lanchester's Equations: [1]

$$\dot{x}_i = -x_i u_i - x_i \sum_j a_{ij} y_j - \sum_j b_{ij} y_j + r_i; \quad i = 1, 2, \dots, I$$

$$\dot{y}_j = -y_j v_j - y_j \sum_i x_i c_{ij} - \sum_i x_i d_{ij} + s_j; \quad j = 1, 2, \dots, J$$

where

the x_i and y_j represent Blue and Orange forces of different types, the u_i and v_j are self-attrition coefficients, the a_{ij} and c_{ij} are area-fire attrition coefficients, the b_{ij} and d_{ij} are aimed-fire attrition coefficients, and the r_i and s_j are replenishment coefficients.

We constrain*

$$x_i, y_j \geq 0 \quad ; \quad \text{for all } i, j$$

and

$$\sum_i x_i = X, \quad \sum_j y_j = Y$$

If any of the force component levels is zero, the corresponding equation is to be dropped out altogether. It

*It is also possible to introduce constraints on the individual types of forces, when more than one combat site is involved. We could also consider the weighted constraint $X = \sum_i h_i x_i$

does not make physical sense to assign attrition against non-existing assets. Eqs. (1), less the linear terms, are known as Lotka-Volterra equations. They have been used extensively to model complex ecosystems in which the stability of a given population distribution is a central question [2].

The question of an optimum force distribution for our system may be approached by choosing an objective function (i.e., a measure of effectiveness) for the combat processes. A useful choice for our purposes is the difference of the weighted sums of the loss rates:

$$M \triangleq \sum_i f_i(\dot{x}_i - r_i) - \sum_j g_j(\dot{y}_j - s_j)$$

For simplicity, we consider here only the case in which all the weighting coefficients are unity, so that

$$M = \sum_i (\dot{x}_i - r_i) - \sum_j (\dot{y}_j - s_j) \quad (2)$$

The first major result of our study to date is that, for this choice of M, there exist optimum force distribution (row and column) vectors x^* and y^* such that for any other vectors x and y

$$x^* A y^* \leq M^* \leq x^* A y \quad (3)$$

where

$$M^* \triangleq x^* A y^*$$

and

\hat{A} is a matrix determined by the attrition coefficients.

This result, derived in Appendix A, is a direct consequence of the fact that the objective function M can be written in a form mathematically equivalent to a matrix game. The optimum vectors x^* and y^* can therefore be determined by solving a linear program. It follows that fairly large problems (say, a few hundreds of variables) can be handled with a reasonable computation load.

A difficulty arises with the problem formulation, however, when the optimal vectors x^* and y^* have zero components, either because the solution is degenerate or because the orders of x and y are different. In such cases, the objective function M is defective, in that it incorporates aimed-fire attrition against non-existent forces. How best then to proceed deserves further study; several alternatives (none obviously superior to the others) come readily to mind.

A second difficulty enters because of the fact that Eq. 1 is non-linear, and can in general only be integrated numerically. In particular, the optimization is valid only at the initial instant, since the force distributions will usually change continuously as the attrition processes evolve in time.

Dynamical Considerations

Variations of force composition can be suppressed by the artifice of choosing the replenishment rates r_i and s_j to make

$$\dot{x}_i = \dot{y}_j = 0; \quad i = 1, 2, \dots, I; \quad j = 1, 2, \dots, J$$

at $x = x^*$ and $y = y^*$.

Since the optimizing vectors x^* and y^* are unaffected by these replenishment terms, the operating point (x^*, y^*) is then in equilibrium and the cost of combat becomes

$$(M = R - S - M^* = S - R)$$

where

$$R \triangleq \sum_i r_i, \quad S \triangleq \sum_j s_j$$

If the combat processes remain in equilibrium for T time units, the integrated cost is

$$TM^* \quad (4)$$

Whether or not the equilibrium can be maintained, of course, depends on whether or not the operating point (x^*, y^*) exhibits neighborhood stability. To investigate this question, we linearize Eq. 1 by substituting

$$x_i \leftarrow x_i^* + \delta x_i, \quad y_j \leftarrow y_j^* + \delta y_j$$

and discarding terms of 2nd order. The result is the linear system

$$\begin{bmatrix} \delta \dot{x} \\ \delta \dot{y} \end{bmatrix} = -\bar{C} \begin{bmatrix} \delta x \\ \delta y \end{bmatrix}$$

in which the "conflict matrix" has the form

$$\bar{C} = \begin{bmatrix} \hat{A} & \hat{B} \\ \hat{D} & \hat{C} \end{bmatrix}$$

The elements of \bar{C} are real and positive, and the submatrices \hat{A} and \hat{C} are diagonal.

It is shown in Appendix B that the matrix \bar{C} exhibits the property that

$$-\hat{a}_{11} + \sum_i \hat{a}_{21} = k_1 \quad ; \quad \text{for all } i$$

and

$$-\hat{c}_{jj} + \sum_k \hat{b}_{kj} = k_2 \quad ; \quad \text{for all } j$$

from which it follows immediately that

$$\delta \dot{x} - k_1 \delta x = \delta \dot{y} - k_2 \delta y \quad (6)$$

where

$$\delta x \triangleq \sum_i \delta x_i \quad ; \quad \delta y \triangleq \sum_j \delta y_j$$

are the perturbations of the aggregated forces around the operating point (x^*, y^*) .

The column-sum parameters k_1 and k_2 tell us a great deal about neighborhood stability. In particular, it is shown in the appendix that the operating point (x^*, y^*) is locally stable if k_1 and k_2 are both negative. If $k_1 = k_2 = k > 0$, then k is the largest eigenvalue of $-\bar{C}$ and the system is unstable. We conjecture, but have not yet proven, that for $0 < k_1 \leq k_2$, $-\bar{C}$ has a largest real eigenvalue λ_{\max} with $k_1 < \lambda_{\max} \leq k_2$. If true, it follows from Eq. 6 that

$$\delta x / \delta y < 0 \quad \text{for } k_1 < \lambda_{\max} < k_2 \quad (7)$$

so that departures of the aggregated forces from an unstable optimum equilibrium will be of opposite sign such that if one force grows, the other must shrink. The situation with k_1 and k_2 of opposite signs is more complicated and requires additional study, but we anticipate that a similar phase relationship will hold for all unstable cases, i.e., whenever the corresponding eigenvalue is positive.

Example

The three aggregate parameters M^* , k_1 and k_2 summarize the interplay within a rich non-linear dynamic system of great complexity. Even the 2×2 case involves 26 parameters, all of which can in general be non-zero. The following numerical example serves to indicate the diversity of behavior patterns which can ensue. For this example, the bilinear attrition coefficients are given by the matrices

$$A = \begin{bmatrix} 1.0 & .3 \\ .6 & .9 \end{bmatrix} \quad C = \begin{bmatrix} 1.2 & 1.0 \\ 1.1 & .6 \end{bmatrix},$$

the linear attrition coefficients by

$$B = \begin{bmatrix} .15 & .1 \\ .15 & .3 \end{bmatrix} \quad D = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix},$$

and the self-attrition coefficients by

$$U = \begin{bmatrix} .3 & .3 \end{bmatrix} \quad V = \begin{bmatrix} .1 \\ .2 \end{bmatrix}$$

The results of optimization are tabulated in Table 1, in which for convenience we have held the XY product constant at 4. We note that the operating point (x^*, y^*) is unstable for small values of X, and stable for large values. When the signs of k_1 and k_2 are the same, λ_{\max} also has that sign. But for mixed signs, λ_{\max} may be + or -. [A conjecture that the largest absolute value governs in such cases can be disproved by counterexample].

Finally, we note that the cost of combat, M^* , favors Orange (e.g., y) for both small and large values of X, whereas Blue is favored in the intermediate cases. Why this happens becomes clear in terms of the individual aggregate loss rates \dot{X} and \dot{Y} . For example, $|\dot{X}|$ is high for $X = 1/2$ because of the adverse force ratio. As X increases, $|\dot{X}|$ decreases because Y decreases. Finally, when $X = 4$ the area fire losses cause $|\dot{X}|$ to increase again. The same factors affect the behavior of $|\dot{Y}|$, so that M^* is the difference of two U-shaped curves. It is nonetheless reassuring that the fractional loss rates, $|\dot{X}/X|$ and $|\dot{Y}/Y|$, behave monotonically, as conventional wisdom would predict.

Figure 1 shows two unstable phase trajectories for the case $X = 1$, $Y = 4$. These were obtained by perturbing the optimum system from its equilibrium point by very small amounts first in one, and then in the opposite, direction along the eigenvector corresponding to the dominant root $\lambda_{\max} = .137$. In the first case, indicated by dots every 40 time units, the system diverges from (x^*, y^*) with X increasing and Y decreasing. However, because of the nonlinearity of the generalized Lanchester equations, there are multiple equilibrium points. In this example, one of these is at $(x_1 = 1.77, x_2 = 1.2862)$ and $(y_1 = 1.067, y_2 = .352)$ corresponding to $X = 3.056$, $Y = 1.419$. Furthermore, this equilibrium point is stable and the system equalizes itself there, far from the original operating point. Here even though Blue gains the initial advantage, he is unable to completely eliminate Orange.

In contrast, when Orange gains the initial advantage, he completely eliminates Blue in less than 120 time steps. This trajectory (the one on the left) was obtained by perturbing the equilibrium solution in the opposite direction along the eigenvector of λ_{\max} .

Discussion and Recommendations

The many issues which remain outstanding divide into two main categories: technical and operational. Questions concerning stability are prominent among the

former, with proof of the conjecture that $k_1, k_2 > 0$ implies instability, and methods for estimating the dominant eigenvalue when k_1 and k_2 are of opposite signs, both having high interest. In addition, it is important to extend the stability analysis to include the case where forces of different types are constrained separately.

How to obtain approximately optimum force distributions when the solution of the linear program is degenerate also deserves careful study; and a great deal of analysis plus numerical experimentation should be devoted to understanding the influence of the various parameters on aggregate parameters such as M^* , k_1 , k_2 , \dot{X}/X and \dot{Y}/Y .

A rich variety of operational questions appear to fall within the purview of our model. By way of illustration, consider the scenario of Figure 2. There are three battle areas, labelled I, II and III, with three different types of Blue forces and two types of Orange. The most obvious question, how should the forces be distributed, given X and Y, has already been discussed. It is worth note in passing, however, that objective functions of the form

$$M = C_1 \frac{\dot{X}}{X} - C_2 \frac{\dot{Y}}{Y}$$

may be particularly relevant to military operations; the coefficients C_1 and C_2 would presumably be chosen to reflect a commander's willingness to accept losses, with the smaller C_1 attaching to the offensive force.

A more subtle question concerns which battle areas to defend (or attack). Given the attrition coefficients and the aggregate force levels, an unattractive area will be assigned zero forces by the optimization routine. In order to "shorten the lines," however, we could ask by how much would the attrition coefficients need to be changed (say, by building field fortifications) before the area would become defensible. Or in which area should the field fortifications be constructed?

Additional questions address the force level necessary to achieve the military objective. For example, how large must Y be (for given X) before the combat becomes unstable (i.e., before a breakthrough can be achieved?) and in which of the three areas will the breakthrough occur. For this we must look to the eigenvector corresponding to the dominant eigenvalue. Conversely, for the defender, how large must X be in order to maintain stability and prevent a breakthrough? It seems certain that the modes of instability, and the terminal states to which they lead, will be of far greater operational importance than the value of M^* itself.

To answer questions such as those above in a meaningful way will require a great deal of additional work. In particular, techniques for estimating and validating attrition coefficients appropriate to a specified style and place of combat need to be developed. The concepts considered here may provide useful guidance in how to aggregate detailed coefficients into higher level combat models; but simulation, controlled experimentation and the analysis of real historical data will be of crucial importance.

Lastly, the theoretical constructs considered thus far have been essentially static, i.e., optimization at one point in time. If the optimum point is unstable, then it seems unreasonable to believe an intelligent commander will not attempt to redistribute his forces in order to capitalize on an incipient breakthrough, or to try and reverse one. Certainly neither will continue to replenish forces at a constant rate, as assumed in our model, but will commit reserves in bulk at an opportune time. When to reoptimize as the force levels diverge from their initial equilibrium remains open to further study. Surely one important question is whether another, stable equilibrium point lies nearby the

unstable optimum point, as in the example presented above; or whether, if left unchecked, a catastrophe is in the making for the side which, perhaps by chance, suffers the first small disadvantage. In addition, if the force levels move too far from their initial values, it may become necessary to modify the system equations to reflect changes in the style and locale of combat.

We conclude by expressing the opinion that an important potential area of application for this and related work lies in war games and simulations aimed at valuating and improving command control and systems and doctrine.

TABLE 1

i	v	x ₁ ⁱ	x ₂ ⁱ	y ₁ ⁱ	y ₂ ⁱ	i	v	x ₁ ⁱ	x ₂ ⁱ	y ₁ ⁱ	y ₂ ⁱ	x ₁ ⁱ	x ₂ ⁱ	y ₁ ⁱ	y ₂ ⁱ
1/2	8	.308	.192	6.192	1.808	-5.842	-5.342	-.900	2.204	.042	.252	11.464	.668		
1	4	.615	.385	3.115	.885	-4.713	-4.855	.142	.942	-.115	.137	4.713	1.214		
2	2	1.231	.769	1.577	.423	-4.340	-4.622	.223	.311	-.431	-.053	2.199	2.311		
4	1	2.46	1.54	.808	.192	-6.848	-4.688	-2.16	-.0039	-1.061	-.719	1.717	4.688		




FIGURE 1: UNSTABLE PHASE TRAJECTORIES




FIGURE 2: MULTIPLE BATTLEFIELDS
The notation x_i means forces of type i assigned to battle-field j .

APPENDIX A: GAME THEORETICAL OPTIMIZATION OF LANCHESTER EQUATIONS

Measure of Effectiveness

Consider the generalized Lanchester Equations

$$\dot{x}_i = -x_i u_i - x_i \sum_j a_{ij} y_j - \sum_j b_{ij} y_j + r_i;$$

$$i = 1, 2, \dots, I$$

$$\dot{y}_j = -y_j v_j - y_j \sum_i c_{ij} x_i - \sum_i d_{ij} x_i + s_j;$$

$$j = 1, 2, \dots, J$$

where the x_i and y_j represent opposing forces of class i and j respectively, and the other parameters are non-negative constants. Furthermore, assume that there are constraints

$$\sum_i x_i = X, \sum_j y_j = Y, x_i \geq 0, y_j \geq 0$$

$$\text{for all } i, j$$

on the forces.

Our objective is to determine advantageous compositions for the forces. To this end, it is convenient to define a measure of effectiveness, M , for the combat by

$$M \triangleq \sum_i f_i (\bar{x}_i - r_i) - \sum_j g_j (\bar{y}_j - s_j)$$

The f_i and g_j are arbitrary weighting factors, which for notational simplicity we hereafter assume to be unity. Then

$$M = \sum_{ij} x_i (c_{ij} - a_{ij}) y_j + \sum_i x_i (d_i - u_i) - \sum_j (b_j - v_j) y_j$$

$$\text{where } d_i \triangleq \sum_j d_{ij}, b_j \triangleq \sum_i b_{ij}$$

This can be rewritten as

$$M = x \tilde{A} y - \tilde{b} y + x \tilde{d} \quad (\text{A.2})$$

where $\tilde{A} \triangleq C - A$: an $I \times J$ rectangular matrix

$\tilde{b} \triangleq b - v$: a J element row vector

$\tilde{d} \triangleq d - u$: an I element column vector

and where x and y are I -element row and J -element column vectors respectively. The expression for M can be written more homogeneously by defining matrices

$$\tilde{B} \triangleq \frac{1}{X} \begin{bmatrix} -\tilde{b} \\ \tilde{b} \\ \vdots \end{bmatrix}, \quad \tilde{D} \triangleq \frac{1}{Y} [\tilde{d}^T; \tilde{d}^T; \dots]$$

so that

$$M = x^* \hat{A} y^* \quad (A.3)$$

$$\text{with } \hat{A} \hat{A} - \hat{B} + \hat{D}$$

Game Formulation

We can interpret M as the payoff of a matrix game by defining

$$p \triangleq x/X, \quad q \triangleq y/Y, \quad \bar{A} \triangleq \hat{A} \hat{A}$$

Then

$$M = p \bar{A} q$$

from which it follows [3] that there are optimum probability vectors p^* and q^* such that, for all x and y ,

$$x \hat{A} y^* \leq M^* \leq x^* \hat{A} y \quad (A.4)$$

$$\text{where } M^* \triangleq x^* \hat{A} y^*$$

$$x^* = X p^*, \quad y^* = Y q^*$$

The vectors p^* and q^* may be found by solving (either of) the dual linear programs

$$\begin{aligned} (P) \min p c & \quad (D) \max b q \\ \text{s.t. } p \bar{A} & \geq b & \quad \text{s.t. } \bar{A} q & \leq c \\ p & \geq 0 & \quad q & \geq 0 \end{aligned}$$

(where b and c are row and column vectors of 1's, having J and I components respectively) and normalizing the solutions by dividing by the (common) optimum value of the objective functions. The reciprocal of this optimum value is M^* .

A result we shall need later involves the complementary slackness conditions: if \bar{p}, q solve the linear program, then

$$\bar{p}_i > 0 \rightarrow \bar{A}_i q = 1; \quad \bar{q}_j > 0 \rightarrow \bar{p} \bar{A}^j = 1$$

with: \bar{A}_i the i th row of \bar{A} and \bar{A}^j the j th column of \bar{A} . It follows that

$$x_i^* > 0 \rightarrow \hat{A}_i y^* = M^*/X \quad (A.5)$$

$$y_j^* > 0 \rightarrow x^* \hat{A}^j = M^*/Y$$

APPENDIX 3: NEIGHBORHOOD STABILITY

The dynamical evolution of Eq. A.1 is complicated by the non-linearity of the equations. Considerable insight may still be gained, however, by selecting (x^*, y^*) as the operating point and choosing the replenishment rates r_i and s_j to make all time derivatives zero. Then (x^*, y^*) is an equilibrium point, and an important question is whether or not the equilibrium is stable.

Aggregate Properties

We consider small perturbations around (x^*, y^*) by letting

$$x_i \leftarrow x_i^* + \delta x_i, \quad y_j \leftarrow y_j^* + \delta y_j$$

for all non-zero components of x^*, y^* . Neglecting 2nd order terms, we then have from Eq. A.2

$$\sum_i \delta x_i - \sum_j \delta y_j = \delta x [\hat{A} y^* + \hat{d}] + [x^* \hat{A} - \hat{b}] \delta y$$

But Eqs. A.5 and A.3 together imply that

$$\hat{A}_i y^* = (\hat{A}_i - \hat{B}_i + \hat{D}_i) y^* = \hat{A}_i y^* - \frac{1}{X} \hat{b}_i y^* + \hat{d}_i$$

so that, for all i ,

$$\hat{A}_i y^* + \hat{d}_i = (M^* + \hat{b}_i y^*)/X \triangleq k_1 \quad (B.1)$$

Similarly, for all j ,

$$x^* \hat{A}^j - \hat{b}_j = (M^* - x^* \hat{d})/Y \triangleq -k_2$$

It follows that

$$\delta \dot{X} - k_1 \delta X = \delta \dot{Y} - k_2 \delta Y \quad (B.2)$$

$$\text{where } \delta X \triangleq \sum_i \delta x_i, \quad \delta Y \triangleq \sum_j \delta y_j$$

The parameters k_1 and k_2 serve to characterize partially the dynamical relationship between the aggregated perturbations δX and δY of the opposing forces.

Stability Analysis

In terms of the perturbation vectors δx and δy , Eq. A.1 can be written as

$$\begin{bmatrix} \delta \dot{x} \\ \delta \dot{y} \end{bmatrix} = - \begin{bmatrix} \hat{A} & \hat{B} \\ \hat{D} & \hat{C} \end{bmatrix} \begin{bmatrix} \delta x \\ \delta y \end{bmatrix} \quad (B.3)$$

$$\begin{aligned} \text{where: } \hat{a}_{ii} &= A_i y^* + u_i & \hat{b}_{ij} &= x_i^* a_{ij} + b_{ij} \\ \hat{d}_{ii} &= c_{ii} y_i^* + d_{ii} & \hat{c}_{jj} &= x^* c^j + v_j \\ \text{and } \hat{a}_{ij} &= \hat{c}_{ij} = 0 & \text{for all } j \neq i \end{aligned}$$

The operating point (x^*, y^*) will be stable iff all eigenvalues of the "conflict matrix"

$$\tilde{C} \triangleq - \begin{bmatrix} \hat{A} & \hat{B} \\ \hat{D} & \hat{C} \end{bmatrix}$$

are in the left half plane. By similarity, an equivalent statement is that all eigenvalues of the matrix

$$\tilde{C} \triangleq \begin{bmatrix} -\hat{A} & \hat{B} \\ \hat{D} & -\hat{C} \end{bmatrix} \quad (B.4)$$

must be in the LHP.

Note that the column sums on the left side of \tilde{C} are given by

$$\sum_i \tilde{C}_{ii} = -(A_i y^* + u_i) + \sum_i (c_{ii} y_i^* + d_{ii}) = k_1$$

and for columns on the right by

$$\sum_j \tilde{C}_{ij} = -(x^* c^j + v_j) + \sum_j (x_i^* a_{ij} + b_{ij}) = k_2$$

We make two important observations: First, all of the Gershgorin disks [4] of \tilde{C} lie totally within the LHP whenever

$$k_1 < 0, \quad k_2 < 0 \quad (B.5)$$

which is therefore a sufficient condition for (x^*, y^*) to be a stable operating point.

Second, if

$$k_1 = k_2 = k > 0$$

(B.6)

then k is the dominant eigenvalue of \tilde{C} and the system is unstable. That k is an eigenvalue of \tilde{C} is seen from the fact that

$$\text{Det}[\tilde{C} - kI] = 0$$

since any row of $\tilde{C} - kI$ can be brought to all zeros by replacing it with the sums of all the rows. Consideration of the Gershgorin disks of \tilde{C} , all of which pass through k , show it to be the eigenvalue with largest real part. We conjecture, but have only been able to prove for special cases, that \tilde{C} will always be unstable when k_1 and k_2 are both positive. The stability of \tilde{C} is not determinable from k_1 and k_2 alone when they are of opposite sign.

REFERENCES

- [1] L. Delansky, "Present State of the Lanchester Theory of Combat," Opns. Res. 12, 344-358 (1964).
- [2] R. M. May, Stability and Complexity in Model Ecosystems, Princeton Univ. Press, 1974, 2nd Edition.
- [3] G. Strang, Linear Algebra and Its Applications, Academic Press, 1976, pp. 319-325.
- [4] T. S. Chow, "Matrices and Linear Algebra," in Handbook of Applied Mathematics, (C. E. Pearson, ed.) p. 906.

INFORMATION PROCESSING AND DECISIONMAKING ORGANIZATIONS: A MATHEMATICAL DESCRIPTION

by

Alexander H. Lewis

Laboratory for Information and Decision Systems
 Massachusetts Institute of Technology
 Cambridge, Massachusetts 02139

ABSTRACT

An overview of an analytical approach to the modeling and evaluation of information processing and decisionmaking organizations is presented. The mathematical framework used in modeling the individual decisionmakers as well as the organization is that of n-dimensional information theory. The data flow formalism is used to model in a precise manner the various types of interactions between decisionmakers as well as interactions between humans and the command, control, and communication system that supports the organization. Comparison and evaluation of alternative organizational forms is accomplished by considering organizational performance, individual workload, and the sets of satisfying decision strategies.

1. INTRODUCTION

In considering organizational structures for teams of decisionmakers, a designer must address the questions of who receives what information and who is assigned to make which decisions. The resolution of these questions specifies the organizational form. The designer's problem is the selection of a form so that the resulting organization meets its performance specifications and the individual members are not overloaded, i.e., the task requirements do not exceed their individual processing limitations.

While the role of the human decisionmakers is central to the design problem, the latter cannot be decoupled from the consideration of the information system that supports the organization. Consider, for example, a tactical military organization supported by a command, control, and communications (C³) system. Information is collected from many sources, distributed to appropriate units in the organization for processing, and used by the commanders and their staff to make decisions. These decisions are then passed to the units responsible for carrying them out. Thus, a given organization design implies the existence of a C³ system that supports it. Conversely, the presence of a C³ system in support of an organization modifies the latter's operations; it may create operational modes not foreseen during the organizational design phase. Therefore, if a quantitative description of the organization design problem is to be developed, it must take into account not only the organization members, but also the collection of equipment and procedures that constitute the organization's C³ system.

In order to develop a quantitative methodology for the analysis and evaluation of information processing and decisionmaking organizations, it is necessary that a set of compatible models be obtained that describe the organization and its environment. This modeling effort has been divided in three steps.

The first one is the modeling of the tasks the organization is to execute and the definition of the boundary between the organization and its environment. The second step is the selection of mathematical models that describe the members of the organization. The third step is the modeling of organizational form, i.e., the specification of the information and decision structures that characterize the organization. This step includes the specification of the protocols for information exchange and the modeling of the communication systems, the data bases, and the decision aids that the organization uses to perform its tasks.

The methodology itself consists of two main parts. In the first one, the analysis of the organization, the models are used to describe the organization in terms of a locus defined on a generalized performance - workload space. This locus is obtained by computing an index of performance for the organization and measures of the workload for each individual member of the organization as functions of the admissible decision strategies used by the decisionmakers. The second part of the methodology addresses the question of evaluating organizational designs and comparing alternative structures.

The analytical framework used for modeling the tasks, the individual organization members, the C³ system, and the organization as a whole is that of n-dimensional information theory [13]. A brief description of the key quantities and of the partition law of information [5] is presented in the next section.

2. INFORMATION THEORETIC FRAMEWORK

Information theory was first developed as an application in communication theory [15]. But, as Khinchin [9] showed, it is also a valid mathematical theory in its own right, and it is useful for applications in many disciplines, including the modeling of simple human decisionmaking processes [16] and the analysis of information-processing systems.

There are two quantities of primary interest in information theory. The first of these is entropy: given a variable x , which is an element of the alphabet X , and occurs with probability $p(x)$, the entropy of x , $H(x)$, is defined to be

$$H(x) = - \sum_x p(x) \log p(x) \quad (2.1)$$

and is measured in bits when the base of the logarithm is two. The other quantity of interest is average mutual information or transmission: given two variables x and y , elements of the alphabets X and Y ,

and given $p(x)$, $p(y)$, and $p(x|y)$ (the conditional probability of x , given the value of y), the transmission between x and y , $T(x:y)$ is defined to be

$$T(x:y) = H(x) - H_y(x) \quad (2.2)$$

where

$$H_y(x) = - \sum_y p(y) \sum_x p(x|y) \log p(x|y) \quad (2.3)$$

is the conditional uncertainty in the variable x , given full knowledge of the value of the variable y .

McGill [13] generalized this basic two-variable input-output theory to N dimensions by extending Eq. (2.2):

$$T(x_1, x_2, \dots, x_N) = \sum_{i=1}^N H(x_i) - H(x_1, x_2, \dots, x_N) \quad (2.4)$$

For the modeling of memory and of sequential inputs which are dependent on each other, the use of the entropy rate, $\bar{H}(x)$, which describes the average entropy of x per unit time, is appropriate:

$$\bar{H}(x) = \lim_{m \rightarrow \infty} \frac{1}{m} H(x(t), x(t+1), \dots, x(t+m-1)) \quad (2.5)$$

Transmission rates, $\bar{T}(x:y)$, are defined exactly like transmission, but using entropy rates in the definition rather than entropies.

The Partition Law of Information [5] is defined for a system with $N-1$ internal variables, w_1 through w_{N-1} , and an output variable, y , also called w_N . The law states

$$\begin{aligned} \sum_{i=1}^N H(w_i) &= T(x:y) + T_y(x; w_1, w_2, \dots, w_{N-1}) \\ &+ T(w_1, w_2, \dots, w_{N-1}; y) \\ &+ H_x(w_1, w_2, \dots, w_{N-1}, y) \end{aligned} \quad (2.6)$$

and is easily derived using information theoretic identities. The left-hand side of (2.6) refers to the total activity of the system, also designated by G . Each of the quantities on the right-hand side has its own interpretation. The first term, $T(x:y)$, is called throughput and is designated G_t . It measures the amount by which the output of the system is related to the input. The second quantity,

$$\begin{aligned} T_y(x; w_1, w_2, \dots, w_{N-1}) &= T(x; w_1, w_2, \dots, w_{N-1}, y) \\ &- T(x:y) \end{aligned} \quad (2.7)$$

is called blockage and is designated G_b . Blockage may be thought of as the amount of information in the

input to the system that is not included in the output. The third term, $T(w_1, w_2, \dots, w_{N-1}; y)$ is called coordination and designated G_c . It is the N -dimensional transmission of the system, i.e., the amount by which all of the internal variables in the system constrain each other. The last term, $H_x(w_1, w_2, \dots, w_{N-1}, y)$, designated by G_n represents the uncertainty that remains in the system variables when the input is completely known. This noise should not be construed to be necessarily undesirable, as it is in communication theory: it may also be thought of as internally-generated information supplied by the system to supplement the input and facilitate the decisionmaking process. The partition law may be abbreviated:

$$G = G_t + G_b + G_c + G_n \quad (2.8)$$

A statement completely analogous to (2.8) can be made about information rates by substituting entropy rate and transmission rates in (2.6).

3. TASK MODEL [8,16]

The organization, perceived as an open system [10], interacts with its environment; it receives signals or messages in various forms that contain information relevant to the organization's tasks. These messages must be identified, analyzed, and transmitted to their appropriate destinations within the organization. From this perspective, the organization acts as an information user.

Let the organization receive data from one or more sources external to it. Every τ_n units of time on the average, each source n generates symbols, signals, or messages x_{ni} from its associated alphabet X_n , with probability p_{ni} , i.e.,

$$p_{ni} = p(x_{ni} = x_{ni}) ; x_{ni} \in X_n \quad i = 1, 2, \dots, \gamma_n \quad (3.1)$$

$$\sum_{i=1}^{\gamma_n} p_{ni} = 1 ; n = 1, 2, \dots, N' \quad (3.2)$$

where γ_n is the dimension of X_n . Therefore, $1/\tau_n$ is the mean frequency of symbol generation from source n .

The organization's task is defined as the processing of the input symbols x_n to produce output symbols. This definition implies that the organization designer knows a priori the set of desired responses Y and, furthermore, has a function or table $L(x_n)$ that associates a desired response or a set of desired responses, elements of Y , to each input $x_n \in X_n$.

It is assumed that a specific complex task that must be performed can be modeled by N' sources of data. Rather than considering these sources separately, one supersource composed of these N' sources is created. The input symbol \underline{x} may be represented by an N' -dimensional vector with each of the sources represented by a component of this vector, i.e.,

$$\underline{x}' = (x_1, x_2, \dots, x_N) \quad ; \quad \underline{x}' \in X \quad (3.3)$$

To determine the probability that symbol \underline{x}' is generated, the independence between components must be considered. If all components are mutually independent, then p_j is the product of the probabilities that each component of \underline{x}' takes on its respective value from its associated alphabet:

$$p_j = \prod_{n=1}^N p_{nj} \quad (3.4)$$

If two or more components are probabilistically dependent on each other, but as a group are mutually independent from all other components of the input vector, then these dependent components can be treated as one supercomponent, with a new alphabet. Then a new input vector, \underline{x} , is defined, composed of the mutually independent components and these super-components.

This model of the sources implies synchronization between the generation of the individual source elements so that they may, in fact, be treated as one input symbol. Specifically, it is assumed that the mean interarrival time for each component τ_n is equal to τ . It is also assumed that the generation of a particular input vector, \underline{x}_j , is independent of the symbols generated prior to or after it.

The last assumption can be weakened, if the source is a discrete stationary ergodic one with constant interarrival time τ that could be approximated by a Markov source. Then the information theoretic framework can be retained [8].

The vector output of the source is partitioned into groups of components that are assigned to different organization members. The j -th partition is denoted by \underline{x}^j and is derived from the corresponding partition matrix π^j which has dimension $n_j \times N$ and rank n_j , i.e.,

$$\underline{x}^j = \pi^j \underline{x}. \quad (3.5)$$

Each column of π^j has at most one non-zero element. The resulting vectors \underline{x}^j may have some, all, or no components in common.

The set of partitioning matrices $(\pi^1, \pi^2, \dots, \pi^M)$ shown in Figure 1 specify the components of the input vector received by each member of the subset of decisionmakers that interact directly with the organization's environment. These assignments can be time invariant or time varying. In the latter case, the partition matrix can be expressed as

$$\pi^j(t) = \begin{cases} \pi_o^j & \text{for } t \in T \\ 0 & \text{for } t \notin T \end{cases} \quad (3.6)$$

The times at which a decisionmaker receives inputs for processing can be obtained either through a deterministic (e.g., periodic) or a stochastic rule. The question of how to select the set of partition matrices, i.e., design the information structure between the environment and the organization, has been addressed by Stabile [17,18].




Figure 1. Information Structures for Organizations

4. MODEL OF THE ORGANIZATION MEMBER [2,3,11]

The complete realization of the model of the decisionmaker (DM) who is interacting with other organization members and with the environment is shown schematically in Figure 2.




Figure 2. The Interacting Decisionmaker with Memory

The DM receives signals $\underline{x} \in X$ from the environment with interarrival time τ . A string of signals may be stored first in a buffer so that they can be processed together in the situation assessment (SA) stage. The SA stage contains algorithms that process the incoming signals to obtain the assessed situation \underline{z} . The SA stage may access the memory or internal data base to obtain a set of values d_0 . The assessed situation \underline{z} may be shared with other organization members; concurrently, the DM may receive the supplementary situation assessment \underline{z}' from other parts of the organization; the two sets \underline{z} and \underline{z}' are combined in the information fusion (IF) processing stage to obtain $\bar{\underline{z}}$. Some of the data (d_1) from the IF process may be stored in memory.

The possibility of receiving commands from other organization members is modeled by the variable \underline{v}' and a command interpretation (CI) stage of processing is necessary to combine the situation assessment $\bar{\underline{z}}$ and \underline{v}' to arrive at the choice $\bar{\underline{v}}$ of the appropriate strategy to use in the response selection (RS) stage. The RS stage contains algorithms that produce outputs \underline{y} in response to the situation assessment $\bar{\underline{z}}$ and the command inputs. The RS stage may access data from or store data in memory [7,8].




Figure 3. Detailed Model of the Interacting Decisionmaker

A more detailed description of the decisionmaker model without buffer or memory is shown in Figure 3. This figure shows the internal structure of the four processing stages: SA, IF, CI, and RS. The situation assessment stage consists of a set of U algorithms (deterministic or not) that are capable of producing some situation assessment z . The choice of algorithms is achieved through specification of the internal variable u in accordance with the situation assessment strategy $p(u)$ or $p(u|z)$, if a decision aid (e.g., a preprocessor) is present. A second internal decision is the selection of the algorithm in the RS stage according to the response selection strategy $p(\bar{v}|\bar{z}, v')$. The two strategies, when taken together, constitute the internal decision strategy of the decisionmaker.

The analytical framework presented in Section 2, when applied to the single interacting decisionmaker with deterministic algorithms in the SA and RS stages, yields the four aggregate quantities that characterize the information processing and decisionmaking activity within the DM [2,11]:

Throughput:

$$G_t = T(x, z', v'; z, y) \quad (4.1)$$

Blockage:

$$G_b = H(x, z', v') - G_t \quad (4.2)$$

Internally generated information:

$$G_n = H(u) - H_z(v) \quad (4.3)$$

Coordination:

$$\begin{aligned} G_c = & \sum_{i=1}^U p_i s_{c_i}^1(p(z)) + \alpha_1 H(p_1) + H(z) \\ & + s_{c_i}^{IF}(p(z, z')) + s_{c_i}^{CI}(p(\bar{z}, v')) \\ & + \sum_{j=1}^V p_j s_{c_j}^1(p(\bar{z}|\bar{v} = j)) + \alpha_j H(p_j) + H(y) \\ & + H(z) + H(\bar{z}) + H(\bar{z}, \bar{v}) + T_z(x': z') \\ & + T_z(x', z': v') \end{aligned} \quad (4.4)$$

The expression for G_n shows that it depends on the two internal strategies $p(u)$ and $p(v|\bar{z})$ even though a command input may exist. This implies that the command input v' modifies the DM's internal decision after $p(v|\bar{z})$ has been determined.

In the expressions defining the system coordination, p_i is the probability that algorithm f_i has been selected for processing the input x and p_j is the probability that algorithm h_j has been selected, i.e., $u = i$ and $\bar{v} = j$. The quantities s_{c_i} represent the internal coordinations of the corresponding algorithms and depend on the distribution of their respective inputs; the quantities α_i, α_j are the number of internal variables of the algorithms f_i and h_j , respectively. Finally, the quantity H is the entropy of a binary random variable:

$$H(p) = -p \log_2 p - (1-p) \log_2 (1-p) \quad (4.5)$$

Equations (4.1) to (4.4) determine the total activity G of the decisionmaker according to the partition law of information (2.6). The activity G can be evaluated alternatively as the sum of the marginal uncertainties of each system variable. For any given internal decision strategy, G and its component parts can be computed.

Since the quantity G may be interpreted as the total information processing activity of the system, it can serve as a measure of the workload of the organization member in carrying out his decisionmaking task.

The qualitative notion that the rationality of a human decisionmaker is not perfect, but is bounded [12], has been modeled as a constraint on the total activity G :

$$G = G_t + G_b + G_n + G_c \leq F \tau_0 \quad (4.6)$$

where τ_0 is the symbol interarrival time and F is the maximum rate of information processing that characterizes a decisionmaker. This constraint implies that the decisionmaker must process his input at a rate that is least equal to the rate with which they arrive. For a detailed discussion of this particular model of bounded rationality, see Boettcher and Lewis [2].

Weakening the assumption that the algorithms are deterministic changes the numerical values of G_n and of the coordination term G_c [4]. If memory is present in the model, then additional terms appear in the expressions for the coordination rate and for the internally generated information rate [7,8].

5. ORGANIZATIONAL FORM

In order to define an organizational structure, the interactions between the human decisionmakers that constitute the organization must be specified. The interactions between DMs and the environment have already been described in Section 3. The internal interactions between DMs consist of receiving inputs from other DM's, sharing situation assessments, receiving command inputs, and producing outputs that are either inputs or commands to other DM's. The detailed specification of the interactions requires the determination of what information is to be passed among individual organization members and the precise sequence of processing events, i.e., the standard operating procedure or communication and execution

protocol of the organization.

Information structures that can be modeled within this analytical framework are those that represent synchronized, acyclical information flows. Since inputs are assumed to arrive at a fixed average rate, the organization is constrained to produce outputs at the same average rate. The overall response is made up, in general, of the responses of several members; therefore, each member is assumed to complete the processing corresponding to a particular input at the same average rate.

Within this overall rate synchronization, however, processing of a specific input symbol or vector takes place in an asynchronous manner. If the requisite inputs for a particular stage of processing are present, then processing can begin without regard to any other stage, which implies that concurrent processing is present. For example, as soon as the organization input arrives and is partitioned through π , processing of x begins to obtain z . The IF stage must wait, however, until both the z and z' values are present. Each stage of processing is thus event-driven; a well-defined sequence of events is therefore an essential element of the model specification.

Aacyclical information structures are those whose directed graphs representing the flows of information do not contain any cycles or loops. This restriction is made to avoid deadlock and circulation of messages within the organization. Deadlock occurs when one DM is waiting for a message from another in order to proceed with his task, while the second one is in turn waiting for an input from the first.

The system theoretic representation of the organizational form is useful for showing the various processing stages or subsystems. For example, in Figure 4, a two person organization is shown in block diagram form in which the second member sends information to the first (x^{21}), who in turn can issue commands to the second DM.

Evaluation of the various information theoretic quantities, including total activity, can be accomplished readily, using the decomposition property of the information theoretic framework [5]. However, the internal information structure of the organization is often ambiguous when represented in block diagram terms. For example, the requirement that both z^1 and z^{11} be present before IP^1 processing can begin is not apparent from Figure 4. An alternate representation is needed which shows explicitly the information structure without compromising the usefulness of the information theoretic decomposition property.

The data-flow schema [1,6] has been developed as a model of information flow for systems with asynchronous, concurrent processing activities. Three basic elements are used in their structure: places, transitions, and directed arcs which connect the two. Places and transitions represent conditions and events, respectively. No event occurs unless the requisite conditions are met, but the occurrence of an event gives rise to new conditions. Tokens are used to mark which conditions are in effect; when all input places to (conditions for) a transition contain a token (are satisfied), then the event can occur, which in turn results in the generation of tokens for output places. Since tokens are carriers of data, each transition is a processor which generates a result from the input data and deposits it on an output token which then moves according to the schema's structure along a directed arc to the next stage of processing.

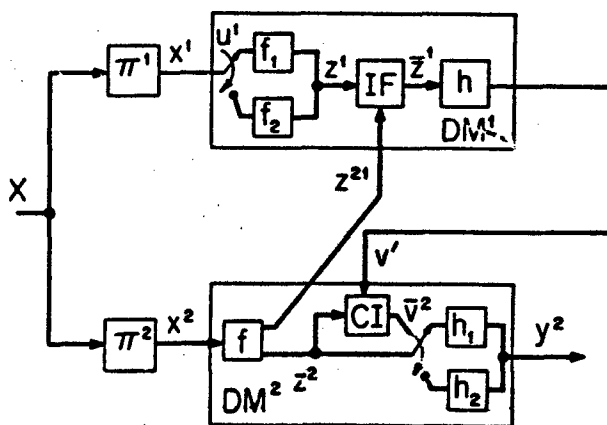


Figure 4. Block diagram representation of two person organization

To represent the information theoretic decisionmaking model using a data-flow formalism, a simple translation in structure is made: distinct inputs and outputs of each subsystem are assigned places and the processing within a subsystem is represented by a transition. Associated with each transition is the set of internal variables of the subsystem, exclusive of the input variables, which are accounted for separately by the input places. By assuming a probability distribution on the organization's inputs, distributions are also included on the places in the structure. Therefore, distributions are also present on subsystem variables, and all information theoretic quantities are well-defined and can be computed as before.

The organization structure shown in Figure 4 can be represented in data-flow terms, as shown in Figure 5. In addition to places, transitions, and directed arcs, the structure contains two new elements, the switches u^1 and \bar{v}^1 . These are logical elements which direct the flow of tokens. The switch u^1 takes values independently, while the value of \bar{v}^1 is determined as a result of the processing by algorithm B^2 contained in CI^2 . Since the structure shown in Figure 5 is equivalent to the system theoretic structure in Figure 4, the internal variable definition and all information theoretic quantities remain unchanged. However, the information structure of the organization is made explicit in Figure 5.

Once an input X is partitioned, the processing by each DM in his respective SA stage (algorithms f) begins concurrently and asynchronously. The information fusion processing (algorithm A^1) must wait until both z^1 and z^{11} have arrived at the input places of IF^1 . Similarly, DM^1 must wait until DM^1 issues a command input v^{13} before the process of command interpretation can begin. This sequence of processing is evident from the representation. Note that because of the assumed synchronization with respect to organization inputs, there can be at most one data token in any single place. The structure is obviously acyclical and deadlock in the organization is prevented.




Figure 5. Data-flow representation of organization structure.

While the data-flow framework provides an equivalent representation for the class of synchronous, acyclical information structure, it is also able to model more general structures, many of which are of interest in the context of organizations. For example, the framework can easily model the cyclic structures which arise when a two-way exchange of information is present in an organization. Such protocols are, of course, common. In addition, fully asynchronous structures can be represented within the framework. Since in a large organization members do not operate at the same rate (same tempo), asynchronous processing is of much interest. The study of these structures and their implications in terms of the n -dimensional information theoretic framework are subjects of current research.

A second advantage of the data flow framework is that it provides a natural way for describing in a precise manner the interactions between the DM's and the data bases and decision aids present in the organization.

The presence of data bases, an integral part of a C^1 system, requires the introduction of two additional modeling elements. The first is the query-response process. The second is the modeling of the data storage devices themselves. Consider, for example, the situation assessment subsystem shown in Figure 6. An accordance with the internal strategy u , an algorithm is chosen to process the input x . However, this algorithm may require parameters (e.g., terrain information, meteorological data) or past situation assessments in order to do the processing. The data base is accessed and queried for this information through the signal D_1 . The data from the data base are provided to the SA subsystem of the DM through D_0 . The same link, D_1 , can be used to update the information in the data base. Clearly, the block diagram representation is ambiguous; the data flow formalism allows for the precise modeling of the fact that data is requested only when certain conditions are met.




Figure 6. Model of SA subsystem with data base access.

Consider next the effect of a data base containing data that do not change during the execution of a task, i.e., the data are fixed. At first glance, it might seem that the addition of the data base with fixed values would have no effect on the total information theoretic rate of activity of the system, i.e.,

$$\bar{H}(d_i) = 0 \quad i = 1, 2, \dots, M \quad (5.1)$$

However, the problem is more complex. For example, if each algorithm f_i accesses β_i parameter values from the data base (in contrast to having these values fixed within the algorithm itself) then the rates of throughput, blockage, and noise of the combined system will not be affected, but the coordination term will have additional activity rate:

$$\Delta \bar{G}_0 = \sum_{i=1}^U \beta_i H[p(u=i)] \quad (5.2)$$

Since a data base increases the overall activity of the system without creating any change in its input-output characteristic, one would question its presence. There are several advantages: (a) reduction in the information that needs to exist within the algorithms or within the decisionmaker model, (b) increased flexibility in the use of algorithms and hence possible reduction in the number of algorithms, and (c) access to common data by several organization members. Even though there is increased coordination activity due to the interaction between the DM and the data base, the total activity of the DM may be reduced — the task may be redesigned to fall within the bounded rationality constraints.

Similar arguments apply to the modeling and analysis of decision aids. Preliminary results indicate that an inappropriately designed decision aid may not reduce a decisionmaker's information processing load, but may actually increase it [4].

In this section, an approach to modeling the organizational form — the specification of the protocols for interaction between DM's — and the supporting command, control, and communication system has been presented. It is based on an integration of the data flow formalism with the information theoretic framework used in the quantitative modeling of the decisionmaking process.

6. ANALYSIS OF ORGANIZATIONS

As stated in Section 3, it is assumed that the designer knows a priori the set of desired responses Y to the input set X . Then the performance of the organization in accomplishing its tasks can be evaluated using the approach shown in Figure 7.




Figure 7. Performance evaluation of an organization

The organization's actual response y can be compared to the desired set Y_d and a cost assigned using a cost function $d(y,Y)$. The expected value of the cost, obtained by averaging over all possible inputs, can serve as a performance index, J , for the organization. For example, if the function $d(y,Y)$ takes the value of zero when the actual response is one of the desired ones and unity otherwise, then

$$J = E \{d(y,Y)\} = p(y \neq Y_d) \quad (6.1)$$

In this case, J represents the probability of the organization making the wrong decision, i.e., the probability of error. Once the organizational form is specified, the total processing activity G and the value of organizational performance J can be expressed as functions of the internal decision strategies selected by each decisionmaker. Let an internal strategy for a given decisionmaker be defined as pure, if both the situation assessment strategy $p(u)$ and the response selection strategy $p(v|Z)$ are pure, i.e., an algorithm f_i is selected with probability one and an algorithm h_j is selected also with probability one when the situation assessed as being Z :

$$D_k = (p(u=i) = 1 ; p(v=j|Z=Z) = 1) \quad (6.2)$$

for some i , some j , and for each Z element of the alphabet Z . There are n possible pure internal strategies,

$$n = U \cdot V^M \quad (6.3)$$

where U is the number of f algorithms in the SA stage, V the number of h algorithm in the RS stage and M the dimension of the set Z . All other internal strategies are mixed [14] and are obtained as convex combinations of pure strategies:

$$D(p_k) = \sum_{k=1}^n p_k D_k \quad (6.4)$$

where the weighting coefficients are probabilities.

Corresponding to each $D(p_k)$ is a point in the simplex

$$\sum_{k=1}^n p_k = 1, \quad p_k \geq 0 \quad \forall k \quad (6.5)$$

The possible strategies for an individual DM are elements of a closed convex hyperpolyhedron of dimension $n-1$ whose vertices are the unit vectors corresponding to pure strategies.

Because of the possible interactions among organization members, the value of G depends not only on $D(p_k)$ but also on the internal decisions of the other decisionmakers. A pure organizational strategy is defined as a M -tuple of pure strategies, one from each DM:

$$A_{1,2,\dots,M} = \{D_{k_1}, D_{k_2}, \dots, D_{k_M}\} \quad (6.6)$$

Independent internal decision strategies for each DM, whether pure or mixed, induce a behavioral strategy [14] for the organization, which can be expressed as

$$A = \sum_{1,2,\dots,M} (A_{1,2,\dots,M} \prod_{i=1}^M p_{k_i}) \quad (6.7)$$

where p_{k_i} is the probability of using pure strategy, D_{k_i} . Because each DM is assumed to select his strategy independently of other DM's, the strategy space of the organization, S^0 , is determined as the direct sum of the individual DM strategy spaces:

$$S^0 = S^1 \oplus S^2 \oplus \dots \oplus S^M \quad (6.8)$$

where S^i denotes the individual DM strategy space. The dimension of S^0 is given by

$$s = \dim S^0 = \sum_{i=1}^M (n_i - 1)$$

Thus, the organizational strategies are elements of an s -dimensional closed convex hyperpolyhedron.

As A ranges over S^0 , the corresponding values of the performance index J and the activity or workload of each individual organization member can be computed. In this manner, the set S^0 is mapped into a locus on the $M+1$ dimensional performance-workload space, namely the space $(J, G^1, G^2, \dots, G^M)$. Note that only the internal processing activity of the decisionmakers is presented in the locus and not the total activity of the system which includes the activity of the decision aids, data bases, and other components of the supporting C^1 system. Consequently, the bounded rationality constraints become hyperplanes in the performance-workload space. Since the bounded rationality constraint for all DM's depends on τ , the admissible internal decision strategies of each DM will also depend on the tempo of operations. The unconstrained case can be thought of as the limiting case when $\tau \rightarrow \infty$.

The methodology for the analysis of organizational structures allows for the formulation and solution of two problems: (a) the determination of the organizational strategies that minimize J and (b) the determination of the set of strategies for which $J \leq \bar{J}$. The first problem is one of optimization while the latter is formulated so as to obtain satisfying strategies with respect to a performance threshold J . The satisfying condition also defines a plane in the performance-workload space that is normal to the J axis and intersects it at \bar{J} . All points on the locus on or below this plane which also satisfy the bounded rationality constraint for each decisionmaker in the organization define the set of satisfying decision strategies. Analytical properties of this locus as well as a computational approach to its efficient construction have been discussed in [2,3,11].

A qualitative evaluation of an organizational structure can be made by comparing the performance-workload locus to the space defined by the satisfying and bounded rationality constraints. In the same manner, alternative organizational structures can be compared by considering their respective loci.

Since individual decisionmakers select their own decision strategies independently of all other organization members, a particular organizational form can yield a broad range of performance as illustrated by the locus in the performance-workload space. The designer must assess, therefore, the likelihood that strategies which lead to satisfying performance will be selected. A possible measure of this mutual consistency between individually selected strategies can be obtained by comparing the locus of the satisfying strategies to the locus of the organization's strategy space S^0 . Let R^i be the subspaces of organization strategies which are feasible with respect to the bounded rationality constraint of each DM, i.e.,

$$R^i = \{A \mid G^i(A) \leq F^i \tau\} \quad (6.9)$$

and let R^J contain the strategies that satisfy the performance threshold \bar{J} :

$$R^J = \{A \mid J(A) \leq \bar{J}\} \quad (6.10)$$

The subspace of satisfying strategies R^0 is given by:

$$R^0 = R^1 \cap R^2 \cap \dots \cap R^M \cap R^J \quad (6.11)$$


The volume of R^0 , denoted by $V(R^0)$ is compared with that of S^0 , $V(S^0)$, to determine the measure of mutual consistency, Q , i.e.,

$$Q = V(R^0)/V(S^0) \quad (6.12)$$

The ratio Q is a monotonic function of \bar{J} and τ with minimum zero and maximum one. A null value for Q implies that no combination of strategies of the individual decisionmakers will satisfy the design specifications, while unity implies that all organizational strategies are feasible, i.e., satisfy the bounded rationality constraints and the performance specifications.

Since Q can be expressed as a function of \bar{J} and τ only, it can be plotted in the three-dimensional space

(Q, \bar{J}, τ) . A typical plot from a three DM example [3] is shown in Figure 8.



8. Mutual consistency measure Q versus \bar{J} and τ .

7. CONCLUSIONS

An analytical approach to modeling organizational structures for teams of decisionmakers supported by command, control, and communication (C³) systems has been described. The integration of n-dimensional information theory with the data flow schema provides tools for describing the activities and interactions within each decisionmaker model, among decisionmakers, and between a decisionmaker and the supporting C³ system. While only synchronous processing with acyclical information structures has been considered in detail, the approach shows promise for the modeling and analysis of asynchronous information processing and decisionmaking. Furthermore, the introduction of memory in the decisionmaker model, and data bases to the organizational structure has broadened the class of organizations and tasks that can be analyzed using this approach.

8. ACKNOWLEDGEMENT

This paper is an overview of a research program that is being carried out with the participation of several students. The contributions of K. L. Boettcher, G. Chyen, S. A. Hall and D. A. Stabile are gratefully acknowledged.

This work was supported by the Office of Naval Research under contract ONR/N00014-83-K-0183 and the Army Research Institute for the Behavioral and Social Sciences under contract MDA903-83-C-0196.

9. REFERENCES

- [1] Arvind, V. Kathail, and K. Pingali, A dataflow architecture with tagged tokens, MIT Laboratory for Computer Science, Paper No. MIT/LCS-TM-174 (1982).
- [2] K. L. Boettcher and A. H. Lewis, Modeling the interacting decisionmaker with bounded rationality, IEEE Trans. on Systems, Man, and Cybernetics SMC-12 (1982) 334-344.
- [3] K. L. Boettcher and A. H. Lewis, Modeling and analysis of teams of interacting decisionmakers with bounded rationality Automatica 19(6) (1983).

- [4] G. Chyén. Information theoretic models of preprocessors, MIT Laboratory for Information and Decision Systems, SM Thesis (in preparation).
- [5] R. C. Conant, Laws of information which govern systems, IEEE Transactions on Systems, Man, and Cybernetics SMC-6 (1976) 240-255.
- [6] J. B. Dennis, J. B. Fosseen, and J. P. Linderman, Data flow schemas, in: G. Goos and J. Hartmanis (Eds.), Lecture Notes in Computer Science, Volume 5, International Symposium on Theoretical Programming (Springer-Verlag, Berlin, 1974).
- [7] S. A. Hall, Information theoretic models of memory and storage, MIT Laboratory for Information and Decision Systems, SM Thesis LIDS-TH-1232 (1982).
- [8] S. A. Hall and A. H. Lewis, Information theoretic models of memory in human decisionmaking models, in: Proceedings of 6th MIT/ONR Workshop on C³ Systems, LIDS-P-1300 (MIT, 1983).
- [9] A. I. Khinchin, Mathematical Foundations of Information Theory (Dover, New York, 1957).
- [10] E. E. Lawler, III and J. G. Rhode, Information and Control in Organizations (Goodyear Publishing Co., Pacific Palisades, Ca, 1976).
- [11] A. H. Lewis, and K. L. Boettcher, Decisionmaking organizations with acyclical information structures, IEEE Transactions on Systems, Man, and Cybernetics, SMC-13, (3) (1983) 384-391.
- [12] J. G. March, Bounded rationality, ambiguity, and the engineering of choice, Bell Journal of Economics, 9 (1978) 587-608.
- [13] W. J. McGill, Multivariable information transmission, Psychometrika 19(2) (1954) 97-116.
- [14] G. Owen, Games Theory (Saunders, Philadelphia, 1968).
- [15] C. E. Shannon and W. Weaver, The Mathematical Theory of Communication (University of Illinois, Urbana, IL, 1949).
- [16] T. B. Sheridan and W. R. Ferrel, Man-Machine Systems, (MIT Press, Cambridge, MA, 1974).
- [17] D. A. Stabile, The design of information structures: Basic allocation strategies for organizations, MIT Laboratory for Information and Decision Systems, SM Thesis LIDS-TH-1098 (1981).
- [18] D. A. Stabile and A. H. Lewis, The design of information structures: Basic allocation strategies for organizations, in: Proceedings of 6th MIT/ONR Workshop on C³ Systems, LIDS-P-1313 (MIT, 1983).

A MODEL FOR ASYNCHRONOUS DISTRIBUTED COMPUTATION *

Dimitri P. Bertsekas
Laboratory for Information and Decision Systems
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Abstract

We present an algorithmic model for distributed computation of fixed points whereby several processors participate simultaneously in the calculations while exchanging information via communication links. We place, essentially no assumptions on the ordering of computation and communication between processors thereby allowing for completely uncoordinated execution. We find that even under these potentially chaotic circumstances it is possible to solve several important classes of problems including the calculation of fixed points of contraction and monotone mappings arising in linear and nonlinear systems of equations, shortest path problems, and dynamic programming.

1. Introduction

There is presently a great deal of interest in distributed implementations of various iterative algorithms whereby the computational load is shared by several processors while coordination is maintained by information exchange via communication links. In most of the work done in this area the starting point is some iterative algorithm which is guaranteed to converge to the correct solution under the usual circumstances of centralized computation in a single processor. The computational load of the typical iteration is then divided in some way between the available processors, and it is assumed that the processors exchange all necessary information regarding the outcomes of the current iteration before a new iteration can begin.

The mode of operation described above may be termed asynchronous in the sense that each processor must complete its assigned portion of an iteration and communicate the results to every other processor before a new iteration can begin. This assumption certainly enhances the orderly operation of the algorithm and greatly simplifies the convergence analysis. On the other hand synchronous distributed algorithms also have some obvious implementation disadvantages such as the need for an algorithm initiation and iteration synchronization protocol. Furthermore the speed of computation is limited to that of the slowest processor. It is thus interesting to consider algorithms that can tolerate a more flexible ordering of computation and communication between processors. Such algorithms have so far found applications in computer communication networks like the ARPANET [1] where processor failures are common and it is quite complicated to maintain synchronization between the nodes of the entire network as they execute real-time network functions such as the routing algorithm.

*This research was conducted at the M.I.T. Laboratory for Information and Decision Systems with partial support provided by the Defense Advanced Projects Agency under Contract No. ONR-N00014-75-C-1183.

Processor network environments for which weakly coordinated distributed computation seems particularly advantageous typically possess one or more of the following characteristics all of which involve occurrence of some type of unpredictable event.

- 1) Computation nodes and communication links are subject to frequent and/or unexpected failures. (For example packet radio networks).
- 2) Computation nodes have different and/or time varying speeds of execution. (For example each processor is assigned to a perhaps time varying number of tasks involving computation loads which are not fixed a priori).
- 3) Computation at various nodes is event driven. (For example in data collection or sensor networks where the timing, and ordering of measurements may not be predictable.).

It is possible to consider various degrees of coordination in different types of distributed algorithms. An interesting question is to determine the minimum degree of coordination needed in a given algorithm in order to obtain the correct solution. To this end we consider an extreme model of uncoordinated distributed algorithms whereby computation and communication are performed at each processor completely independently of the progress in other processors. It is perhaps surprising that even under these chaotic circumstances it is still possible to solve correctly important classes of fixed point problems. The complete analysis is given in [2] for broad classes of dynamic programming and in [3] for more general fixed point problems involving contraction and monotonicity assumptions. Further related work is [5] and [6].

2. A Model For Distributed Uncoordinated Fixed Point Algorithms

The fixed point problem considered in this paper is defined in terms of a set X , a class F of functions mapping into the extended real line $[-\infty, \infty]$, and a mapping T which maps F into itself. We wish to find an element J^* of F such that

$$J^* = T(J^*) \quad (1)$$

or equivalently

$$J^*(x) = T(J^*)(x), \quad \forall x \in X, \quad (2)$$

where $J^*(x)$ and $T(J^*)(x)$ denote the values of the functions J^* and $T(J^*)$ respectively at the typical element $x \in X$. We will assume throughout that T has a unique fixed point J^* within the set F .

We provide some examples:

Example 1: (Fixed points of mappings on R^n). Let X be the finite set

$$X = \{1, 2, \dots, n\},$$

and F be the set of all real-valued functions on X .

Then F can be identified with the n -dimensional space R^n in the sense that with each $J \in F$ we can associate the n -dimensional vector $J(1), J(2), \dots, J(n)$. Similarly $T(J)$ can be identified with the n -dimensional vector $T(J)(1), \dots, T(J)(n)$, so the fixed point problem (1) amounts to solving the system of n equations

$$J^* = T(J^*) \quad \text{or} \quad J^*(i) = T(J^*)(i), \quad \forall i = 1, \dots, n \quad (3)$$

with the n unknowns $J^*(1), \dots, J^*(n)$. It is also evident that any system of n (possibly nonlinear) equations with n unknowns can be formulated into a fixed point problem such as (3).

Example 2: (Shortest path problems). Let (N, L) be a directed graph where $N = \{1, 2, \dots, n\}$ denotes the set of nodes and L denotes the set of links. Let $N(i)$ denote the downstream neighbors of node i , i.e., the set of nodes j for which (i, j) is a link. Assume that each link (i, j) is assigned a positive scalar a_{ij} referred to as its length. Assume also that there is a directed path to node 1 from every other node. Then it is known ([4], p. 67) that the shortest path distances $J^*(i)$ to node from all other nodes i solve uniquely the equations.

$$J^*(i) = \min_{j \in N(i)} \{a_{ij} + J^*(j)\}, \quad i \neq 1 \quad (4a)$$

$$J^*(1) = 0 \quad (4b)$$

If we make the identifications $X = \{1, 2, \dots, n\}$, F : Set of all functions mapping X into $[0, \infty]$, and define $T(J)$ for all $J \in F$ by means of

$$T(J)(i) = \begin{cases} \min_{j \in N(i)} \{a_{ij} + J(j)\} & \text{if } i \neq 1 \\ 0 & \text{if } i = 1 \end{cases} \quad (5)$$

then we find that the fixed point problem (2) reduces to the shortest path problem.

The shortest path problem above is representative of a broad class of dynamic programming problems which can be viewed as special cases of the fixed point problem (2) and can be correctly solved by using the distributed algorithms of this paper (see [3]).

Our algorithmic model can be described in terms of a collection of n computation centers (or processors) referred to as nodes and denoted $1, 2, \dots, n$. The set X is partitioned into n disjoint sets denoted X_1, \dots, X_n , i.e.

$$X = \bigcup_{i=1}^n X_i, \quad X_i \cap X_j = \emptyset, \quad \text{if } i \neq j.$$

Each node i is assigned the responsibility of computing the values of the solution function J^* [c.f. (1), (2)] at all $x \in X_i$.

At each time instant, node i can be in one of three possible states *compute*, *transmit*, or *idle*. In the *compute* state node i computes a new estimate of the values of the solution function J^* for all $x \in X_i$. In the *transmit* state node i communicates the estimate

obtained from the latest computation to one or more nodes j ($j \neq i$). In the *idle* state node i does nothing related to the solution of the problem. It is assumed that a node can receive a transmission from other nodes simultaneously with computing or transmitting.

We assume that computation and transmission for each node takes place in uninterrupted time intervals $[t_1, t_2]$ with $t_1 < t_2$, but do not exclude the possibility

that a node may be simultaneously transmitting to more than one nodes nor do we assume that the transmission intervals to these nodes have the same origin and/or termination. We also make no assumptions on the length, timing and sequencing of computation and transmission intervals other than the following:

Assumption (A): There exists a positive scalar P such that, for every node i , every time interval of length P contains at least one computation interval for i and at least one transmission interval from i to each node $j \neq i$.

Each node i also has a buffer B_{ij} for each $j \neq i$ where it stores the latest transmission from j , as well as a buffer B_{ii} where it stores its own estimate of values of the solution function for all $x \in X_i$. The contents of each buffer B_{ij} at time t are denoted J_{ij}^t . Thus J_{ij}^t is, for every t , a function from X_j into $[-\infty, \infty]$ and may be viewed as the estimate by node i of the restriction of the solution function J^* on X_j available at time t . The rules according to which the functions J_{ij}^t are updated are as follows:

1) If $[t_1, t_2]$ is a transmission interval from node j to node i the contents $J_{jj}^{t_1}$ of the buffer B_{jj} at time t_1 are transmitted and entered in the buffer B_{ij} at time t_2 , i.e.

$$J_{ij}^{t_2} = J_{jj}^{t_1}. \quad (6)$$

2) If $[t_1, t_2]$ is a computation interval for node i the contents of buffer B_{ii} at time t_2 are replaced by the restriction of the function $T(J_i^{t_1})$ on X_i where, for all t , J_i^t is defined by

$$J_i^t(x) = \begin{cases} J_{ii}^t(x) & \text{if } x \in X_i \\ J_{ij}^t(x) & \text{if } x \in X_j, \quad j \neq i \end{cases} \quad (7)$$

In other words we have

$$J_{ii}^{t_2}(x) = T(J_i^{t_1})(x), \quad \forall x \in X_i \quad (8)$$

3) The contents of a buffer B_{ii} can change only at the end of a computation interval for node i . The contents of a buffer B_{ij} , $j \neq i$ can change only at the end of a transmission interval from j to i .

Additional conditions under which they hold

$$\lim_{t \rightarrow \infty} J_i^t(x) = J^*(x), \quad \forall x \in X_i, i = 1, \dots, n \quad (9)$$

may be found in [2], [3]. An interesting aspect of results of this type is that they do not require that the initial processor buffer contents be identical and indeed these initial conditions can vary within a broad range. This means that for problems that are being solved continuously in real time it is not necessary to reset the initial conditions and resynchronize the algorithm each time the problem data changes. As a result the potential for tracking slow variations on the solution function is improved and algorithms implementation is considerably simplified.

References

- [1] J. McQuillan, G. Falk, and I. Richer, "A Review of the Development and Performance of the ARPANET Routing Algorithm", IEEE Trans. on Communications, Vol. COM-26, 1978, pp. 1802-1811.
- [2] D. P. Bertsekas, "Distributed Dynamic Programming" IEEE Trans. on Automatic Control, Vol. AC-27, 1982, pp. 610-616.
- [3] D. P. Bertsekas, "Distributed Computation of Fixed Points", Report LIDS-P-1135, Laboratory for Information and Decision Systems, Mass. Institute of Technology, Cambridge, Aug., 1981 (to appear in Math. Programming).
- [4] E. L. Lawler, Combinatorial Optimization: Networks and Matroids, Holt, Rinehart, and Winston, N.Y., 1976.
- [5] J. N. Tsitsiklis, D. P. Bertsekas, and M. Athans, "Distributed Deterministic and Stochastic Optimization Algorithms with Applications in System Identification", Proc. of Melecon'83, Athens, Greece, May 1983, paper C5.03.
- [6] J. N. Tsitsiklis, "Problems in Distributed Decision Making", Ph.D. Thesis, Dept. of Electrical Engineering and Computer Science, Mass. Institute of Technology, forthcoming.

A UNIFIED APPROACH TO MODELING AND COMBINING OF EVIDENCE THROUGH RANDOM SET THEORY

I.R. Goodman

Surveillance Systems Department, Code 7223,
Naval Ocean Systems Center, San Diego, California 92152.

Abstract. It has been shown in previous work that generalized fuzzy set theory and infinite-valued logic provide a systematic approach to the modeling and use of both natural language and numerical/statistical information which occurs in the tracking-data association and related problems. This paper continues efforts in establishing connections between these disciplines and classical probability theory. It has been shown that over discrete spaces, probabilistic concepts are all special cases of generalized fuzzy set ones. Conversely, many fuzzy set systems can be shown to be natural extensions of ordinary set operators through isomorphic-like relations with corresponding random set representations via one point coverage functions. Among the new results presented here, it is shown that any fuzzy set membership function has naturally compatible random set and random variable representations. In the latter case, the membership function is the same as the evaluation function of the (non-unique) corresponding random variable over a suitably chosen collection of compound- and, in general, overlapping- sets or events. An application to the classification problem is presented.

INTRODUCTION

Although some investigators in the field of cognitive rationality have come to the conclusion that rational decision making by humans is unobtainable [1], it is the optimistic belief of many others (see, e.g., Cohen [2]) that intuition, properly systematized, can serve as a basis for the choice of a normative theory of decision making. In turn, the latter depends critically upon which fundamental theory or theories of uncertainties and beliefs to accept. Fine [3] considered various approaches involving classical and subjective probability theory with underlying emphasis on modeling random variables rather than random sets. More recently, Freeling [4] has analyzed some nonstandard models, including upper and lower probabilities and the Dempster-Shafer theory of belief, Zadeh's possibility theory, Cohen's inductive probabilities and second order Bayesian probabilities, as well as Shackle's degree of surprise approach. However, without a single basis which can be used for comparisons, Freeling has not been able to demonstrate any deep structural relationships among the various schools of thought, and consequently, comparisons are limited quantitatively. Manes [5] has proposed a formal axiom system for modeling uncertainties which generalizes Zadeh's original fuzzy set theory and is based on the Kleisli category theorem for monoids. (However, it does not capture a large class of generalized fuzzy set systems-see [6].) In addition, Manes' system extends Dempster-Shafer, topological neighborhood, credibility, and other theories. But again, it must be emphasized that this general theory is considering formal similarities of various uncertainty approaches, not their internal structures and their relationships to each other. In addition, there is the ongoing controversies between the proponents of the various approaches viewed somewhat parochially from their respective stands. For example, see Goodman [7] for a listing of arguments between fuzzy set adherents and Bayesian oriented

individuals. Or, see the Lindley "paradox" discussions between Dempster-Shafer and Bayesian backers [8]. Polemics aside, techniques will have to be developed which in some reasonable sense model uncertainties as faithfully as possible so that applications to decision making may be carried out on both the basic intuitive level and normative rigorous level, as discussed by Lindley and Brown in the reconciliation of incoherent data [9].

Random set theory and its modifications (such as through equivalence classes of random sets) could well provide a key to a meaningful analysis of the myriad approaches to the modeling of uncertainties by the establishment of mathematically rigorous relationships between the theory of perception, natural language descriptions, multivalued truth and set theory, and the many schools of thought. The intuitive basis for the use of random sets appears only slightly more complicated than that for ordinary random variables. Indeed, as the new results of this paper show, there is a natural relationship between random set and random variable representations of fuzzy set membership functions.

SUMMARY OF PREVIOUS RESULTS

In this section a brief description of previous results relating random set theory and fuzzy set theory as well as other areas is presented. For basic definitions and background, see [10], [11], [12], [7], [13]

Theorem 1. [7], [10]

Any fuzzy subset A of ordinary set X has, in general, many one point coverage equivalent random set representations $S(A)$, where $S(A)$ is a random subset of X . That is

$$\varphi_A(x) = \Pr(x \in S(A)) \quad ; \quad \forall x \in X. \quad (1)$$

Two examples of this include the nested random set $S_U(A)$ and the very broken-up random set $T(A)$, where

$\phi_A: X \rightarrow [0,1]$ is the membership or possibility function for A and

$$S_U(A) \triangleq \phi_A^{-1}[U,1] \quad (2)$$

while the random membership function for T(A) is

$$\phi_{T(A)} = (\phi_{T(A)}(x))_{x \in X} \quad (3)$$

where each $\phi_{T(A)}(x)$ is a statistically independent zero-one random variable with, for all $x \in X$

$$\Pr(\phi_{T(A)}(x)=1) = \phi_A(x); \Pr(\phi_{T(A)}(x)=0) = 1 - \phi_A(x). \quad (4)$$

Conversely, any random subset S of X has for its one point coverage function $\Pr(x \in S)$ as a function of x, a possibility function for some corresponding unique fuzzy subset A of X.

Thus the fuzzy subsets of a given space X partition it into disjoint and exhaustive parts, each part corresponding to all random subsets of X having a common one point coverage function: the membership function of any fixed fuzzy set.

Theorem 2. [10]

S_U and T as mappings from the collection of all fuzzy subsets of a given space into the collection of all random subsets of the same space induce isomorphic-like relations between various fuzzy set operators and corresponding random set operators—i.e., ordinary set operators on random subsets of the same space. For example, where \approx indicates one point coverage equivalent:

$$f(S_U(A)) \approx S_U(f(A)) \approx f(A) \quad (5)$$

$$S_U(A) \cup S_U(B) \approx S_U(A \cup B) \approx A \cup B \quad (6)$$

$$S_U(A) \cap S_U(B) \approx S_U(A \cap B) \approx A \cap B \quad (7)$$

$$X - S_U(A) \approx S_U(X - A) \approx X - A \quad (8)$$

$$\text{proj}(S_U(A)) \approx S_U(\text{proj}(A)) \approx \text{proj}(A), \quad (9)$$

all fuzzy subsets A, B of space.

Indeed, S_U possesses even stronger properties in that the left-sided equivalences can be replaced by actual equalities for eqs. (5), (6), (7), and (9). Similar results hold for T, but without any of the stronger replacements of one point coverage equivalence by equality. In the case of S_U , the fuzzy set operators for functional transform $f: X \rightarrow Y$, unions \cup , intersections \cap , complements $X - \cdot$, and projections, involve the fuzzy set system $F = (\phi, \psi, \phi_{or}) = (1 - (\cdot), \min, \max)$, for complementation, intersection, and union. In the case of T, the fuzzy set system $F_1 = (1 - (\cdot), \text{prod}, \text{probsum})$ is used for all definitions.

More generally, fuzzy set system $F = (\phi, \psi, \phi_{or})$ consists of a triple of operators, the first being an involution, the second, a t-norm, and the third, a t-conorm [10], [14]. Based on infinitely-valued set theory, ordinary set and logical concepts valid in classical set theory and classical logic may be extended to a fuzzy set form, for any choice of fuzzy set system. For example, consider as above $f: X \rightarrow Y$ and any $y \in Y$ and any fuzzy subset A of X:

$$\begin{aligned} \phi_f(A)(y) &= \text{truth}(y \in f(A)) \\ &= \text{truth}(\exists x (x \in A) \& (y = f(x))) \\ &= \text{truth}(\bigvee_{x \in X} (\phi_A(x) \& (y = f(x)))) \\ &= \text{truth}(\bigvee_{x \in X} (\phi_A(x) \cap f^{-1}(y))) \\ &= \phi_{or}(\text{truth}(x \in A \cap f^{-1}(y))) \\ &= \phi_{or}(\phi_A \cap f^{-1}(y)) \end{aligned}$$

$$\begin{aligned} &= \phi_{or}(\text{truth}((x \in A) \& (y = f^{-1}(y)))) \\ &= \phi_{or}(\phi_A(\text{truth}(x \in A), \text{truth}(y = f^{-1}(y)))) \\ &= \phi_{or}(\phi_A(x), \phi_{f^{-1}}(y)) \\ &= \phi_{or}(\phi_A \cap f^{-1}(y)) \\ &= \phi_{or}(\phi_A(x)) \\ &= \phi_{f^{-1}}(y) \end{aligned} \quad (10)$$

Note that when $F = F'$ (see Theorems 5(a) and 6(b)), the above result reduces to the classical form for the transformation of probability under f. (See also [13].)

Extending the concept of a single one point equivalence map—as are S_U and T, among others—is that of a choice function family. This is a collection of identical maps S_f from the collection of all fuzzy subsets of any space to that of all random subsets of the same space, such that for any spaces X_1, \dots, X_n , and any fuzzy subsets A_1, \dots, A_n of X_1, \dots, X_n , respectively, the corresponding one point equivalent random subsets $S_{f_1}(A_1), \dots, S_{f_n}(A_n)$ have a well-defined distribution, or equivalently, form a well-defined joint collection of random sets. Two important such families are: $S_U \triangleq (S_{U_j})_{j=1,2,\dots}$, where U_1, U_2, \dots form a J-copula stochastic process [10] and $S_{\&}$, formed from repetitive applications of $\psi_{\&}$, a semi-distributive t-norm [10].

Theorem 3. [10], [13], [15]

Theorem 2 can be extended in a natural way for the two choice function families described above.

However, it can be shown that there exist fuzzy set operations which have no random set representation through any possible choice function family. (See [13].) Nevertheless, all of the basic fuzzy set concepts and their generalizations, including systems F_0 , F_1 , and many others, have complete random set representations. (Again, see [13].) If a fuzzy set operation has a random set representation, then that representation must be the restriction of the fuzzy set operation to ordinary sets.

Theorem 4. [10], [13], [15]

Given any choice function family and any ordinary (n-ary) set operation, there exists a unique extension to fuzzy sets such that a random set representation exists relative to that choice function family. In particular, the application to binary compositional set operations generates an easily computable class of extensions.

One of the chief difficulties in carrying out random set representations and interpretations of fuzzy sets and their operations is the non-uniqueness of correspondence. In general, infinitely many random sets can (one point coverage equivalence) represent a given fuzzy set. One method of selecting a single random set representation is through the measure of entropy:

Theorem 5. [16]

Consider any fixed discrete space X. Then

(a) The collection of all random subsets S(A) representing A can be characterized as being in a bijective relation with a simply describable convex linearly bounded subspace R(A) of R^C where $c \triangleq 2^{\text{card}(X)} - 1 - \text{card}(X)$.

(b) The maximal entropy random subset representing A is T(A).

(c) The minimal entropy random subset representing A occurs in a collection of random subsets which is bijective with the vertex set $V(A)$ of $R(A)$, a relatively sparse set.

(d) There is only one nested random subset of the space representing A , namely $S_U(A)$. (This result is true for any space X , discrete or not.) $S_U(A)$ always lies in the class bijective to $V(A)$, although it is not always necessarily the minimal entropy representation.

(e) Note first, that all probability functions and all cumulative and reverse-cumulative probability distribution functions, for ordinary and deficient probability measures, are fuzzy set membership functions. Let A be any fuzzy subset of X with membership function ϕ_A , which is a probability function. Then, that random subset of X which is the singleton-valued one formed from a random variable having ϕ_A as its probability function, represents A and lies in the class bijective to $V(A)$, although as in (d), it does not always possess the minimal entropy.

Some miscellaneous results are summarized next.

Theorem 6.

(a) A general fuzzy set version of the Law of Large Numbers has been established [17]. (For a fuzzy set version of the Central Limit Theorem, see [18].)

(b) Conditional fuzzy sets may be defined analogous to those in classical probability spaces over discrete spaces [17], in turn, yielding a fuzzy set form of Bayes' theorem for general fuzzy set systems F [13], [17]. This may be used to develop, from first principles, a decision theory based upon fuzzy set concepts. In particular, fuzzy set membership functions may be used to model error distributions of attributes and inference rules connecting groups of attributes with parameters of interest. The rules can include statistical tests of hypotheses converted into fuzzy set form by considering their random significance levels as fuzzy set membership functions of the test statistic values. This information may then be combined with observed data to yield a posterior fuzzy set describing the parameters of interest. (See [20], [21] for applications to the target data association problem-PACT algorithm.) When all fuzzy sets involved in the decision problem are also probability functions and the system $F=F'$ (1-., prod, bndsum) is chosen, then all of the above results reduce to ordinary probabilistic ones.

(c) Characterizations have been obtained for random interval representations of fuzzy sets [16].

(d) The Dempster-Shafer theory has been shown to be completely analyzable through random set theory; in particular, by use of Choquet's Capacity theorem concerning subset, superset, and incidence functions of random sets as extensions of the usual properties of probability measures in expanding in alternating sums, unions of events in terms of intersections, and vice-versa. (See [11], [7].)

(e) Although Eytan's claim that Zadeh's version of fuzzy set theory, considered as a formal category, has been pointed out by Pitts, Carrega, and Ponasse not to be a topos (see [22], [23], [24]), Higgs's topos (as discussed in [25]) promises to be a fruitful medium for extending the concept of membership in a fuzzy set to simultaneous membership of several points (in a Higgs set) since a sound and complete intuitionistic logic can be realized through the topos; it is a natural extension of the concept of a set, and indeed a fuzzy set (for single point membership) and is conducive for axiomizations (although see Goguen's earlier axiomization of Zadeh's fuzzy set theory as a category [26]); and finally, it can be shown to be the quotient object

completion and union of a class of natural imbeddings of various forms of Zadeh's fuzzy set theory for multiple memberships [22], [27]. Apropos to multiple point membership of fuzzy sets, Theorem 5 (a) has been extended to the characterizations of all random subsets of X which are multiple point coverage equivalent to a given multiple point membership function, up to any prescribed multiplicity level [16].

Finally, we mention the fundamental result previously obtained concerning how evidence should be combined.

Theorem 7. [19], [20]; [16], for asymptotic properties.

Given any collection of fuzzy subsets A_1, \dots, A_n of a space X describing some common unknown parameter vector, for each choice of non-decreasing combining function $g: [0,1]^n \rightarrow [0,1]$ with respect to any confidence levels of the corresponding level set forms, there is a unique single fuzzy set description of the parameter which minimally contains the information given by A_1, \dots, A_n . This fuzzy set A is determined by the equation

$$\phi_A(x) \triangleq g(\phi_{A_1}(x), \dots, \phi_{A_n}(x)); \text{ all } x \in X. \quad (11)$$

By noting that any t-norm satisfies the conditions for g in the above theorem, an application may be made to a large class of data association problems, where it can be shown that the fuzzy Bayesian solutions of the problems coincide with the combination of evidence approach determined by eq. (11). (See the remarks in Theorem 6(b) concerning the PACT algorithm.)

FUZZY SET MEMBERSHIP FUNCTIONS AS EVALUATIONS OF RANDOM VARIABLES OVER COMPOUND EVENTS

The results of the previous section point out that all fuzzy sets have (in general, many) random set one point coverage representations and that the most common fuzzy set operators also have isomorphic-like random set operator correspondences (under the one point coverage relations). In this section, we present a dual result with respect to random variables over some initial elementary event space. It is first shown (Theorem 8) that given a random variable V over a space X and given any collection A of compound events from X , i.e., A is a collection of subsets of X lying in the σ -algebra on X , a random subset S of A may be constructed such that not only is the one point coverage function of S the same as the evaluation function for V over A , but also that the structure of S is natural with respect to V , i.e., any outcome of S is the filter class of A over a corresponding outcome of V . The next result (Theorem 9) shows the converse: given any random subset S of any given collection A , an elementary event space X for A may be constructed as well as a random variable V over X such that the situation in Theorem 8 holds. The construction, in general is non-unique. Finally, these results are combined with Theorem 1 to show that any fuzzy subset of a space may be represented dually (in many ways, in general, depending on the choice of S) by both the evaluation function of a random variable over compound events and as the one point coverage function of a random subset of the same space, with the random variable and random set in a natural relationship, as mentioned above.

First a basic lemma and definitions are presented.

Lemma and Definitions.

Let X be any space and $\mathcal{A} \subseteq \mathcal{P}(X)$, the class of all ordinary subsets of X . Let $V: \Omega \rightarrow X$ be any random variable relative to probability space $(\Omega, \mathcal{C}, \Pr)$ and some measurable space (X, \mathcal{B}) inducing probability space $(X, \mathcal{B}, \Pr \circ V^{-1})$ with σ -algebra $\mathcal{B} \subseteq \mathcal{A}$.

Let $S: \Omega \rightarrow \mathcal{P}(A)$ be any random subset of A relative to measurable space $(\text{rng}(S), \sigma(S))$, $\text{rng}(S) \subseteq \mathcal{P}(A)$, generated by

$$G(S) \triangleq \{C_{\{a\}}(\text{rng}(S)) \mid a \in A\} \quad (12)$$

where the filter class of $\text{rng}(S)$ on point a is given by

$$C_{\{a\}}(\text{rng}(S)) \triangleq \{B \mid a \in B \in \text{rng}(S)\}, \quad (13)$$

for any $a \in A$. More generally, the filter class of $\text{rng}(S)$ over any $C \in \mathcal{P}(X)$ is given by

$$C_C(\text{rng}(S)) \triangleq \{B \mid C \subseteq B \in \text{rng}(S)\}. \quad (14)$$

(See [13] and [28] for background on random sets.)

Thus S induces probability space $(\text{rng}(S), \sigma(S), \text{Pr} \circ S^{-1})$. Then:

$$V^{-1}(a) = S^{-1}(C_{\{a\}}(\text{rng}(S))), \text{ all } a \in A \quad (15)$$

iff

$$S(\omega) = C_{\{V(\omega)\}}(A), \text{ all } \omega \in \Omega, \quad (16)$$

which mutually imply

$$\text{Pr}(V^{-1}(a)) = \text{Pr}(S^{-1}(C_{\{a\}}(\text{rng}(S)))), \quad (17)$$

i.e.,

$$\text{Pr}(V \in a) = \text{Pr}(V \in S), \text{ all } a \in A. \quad (18)$$

In addition, for any collection of sets such as $G(S)$, a basis $\hat{G}(S)$ may be formed as follows:

(1) Let $A' \subseteq A$ be such that $C_{\{a\}}(\text{rng}(S))$ is a one-to-one function of $a \in A'$, yet

$$G(S) = \{C_{\{a\}}(\text{rng}(S)) \mid a \in A'\}. \quad (19)$$

(2) For any $B \in \mathcal{P}(A')$, define

$$B_B \triangleq C_B(\text{rng}(S)) \rightarrow \bigcup_{a \in A' \cap B} C_{\{a\}}(\text{rng}(S)). \quad (20)$$

(3) Let

$$\hat{G}(S) \triangleq \{B_B \mid B \in \mathcal{H}\}, \quad (21)$$

be the collection of all non-vacuous distinct (and hence disjoint) B_B 's, noting that $\mathcal{H} \subseteq \mathcal{P}(A)$ may be non-unique.

(4) It follows that $\hat{G}(S)$ forms a (disjoint and exhaustive) partitioning of $\bigcup \hat{G}(S)$, where for any $a \in A'$,

$$C_{\{a\}}(\text{rng}(S)) = \bigcup_{a \in B \in \mathcal{H}} B_B \text{ (disjoint)}. \quad (22)$$

(5) In turn, this implies that $\{S^{-1}(B_B) \mid B \in \mathcal{H}\}$ is a partitioning of Ω , where

$$S^{-1}(B_B) \triangleq \{S^{-1}(C) \mid C \in B_B\}$$

and

$$S^{-1}(C) \triangleq \{\omega \mid \omega \in \Omega \ \& \ S(\omega) = C\},$$

by the usual inverse functional notation.

Theorem 8. Random variables generate naturally corresponding random sets whose one point coverage functions match the evaluation functions of the random variables over a given set of compound events.

Let $V: \Omega \rightarrow X$ be a random variable with A arbitrary, where $A \subseteq B \subseteq \mathcal{P}(X)$ as in the lemma. Then define $S: \Omega \rightarrow \mathcal{P}(A)$ by eq. (16). The lemma implies eqs. (15)-(18) hold.

Theorem 9. Converse of Theorem 8: Any random set generates a naturally corresponding random variable and an elementary event space such that the random variable evaluated over certain of the compound events matches the one point coverage function of the random set.

Let only A be given with some random subset $S: \Omega \rightarrow \mathcal{P}(A)$. Note the validity of the development in eqs. (19)-(22). Then there exists a space X (corresponding to elementary events), a random variable $V: \Omega \rightarrow X$ surjective, and a mapping $\Lambda: A \rightarrow \mathcal{P}(X)$ which is injective over A' such that

$$V^{-1}(\hat{a}) = S^{-1}(C_{\{a\}}(\text{rng}(S))), \text{ all } a \in A \quad (23)$$

and equivalently,

$$\hat{S}(\omega) = C_{\{V(\omega)\}}(\hat{A}), \text{ all } \omega \in \Omega, \quad (24)$$

where as usual $\hat{A} \triangleq \{\hat{a} \mid a \in A\}$ in functional form,

and similarly for \hat{S} , yielding immediately

$$\text{Pr}(V \in \hat{a}) = \text{Pr}(a \in S), \text{ all } a \in A. \quad (25)$$

In addition:

(i) Λ must always satisfy-and therefore may be defined by, once V and X are determined-

$$\hat{a} = V(S^{-1}(C_{\{a\}}(\text{rng}(S)))), \text{ all } a \in A. \quad (26)$$

(ii) One choice of V and X is:

$$V \triangleq S, \ X \triangleq \text{rng}(S), \quad (27)$$

in which case Λ reduces to

$$\hat{a} = C_{\{a\}}(\text{rng}(S)), \text{ all } a \in A. \quad (28)$$

(iii) The general solution for V and X (and hence Λ via eq. (26)) is constructed as follows:

From the lemma (5) for each $B \in \mathcal{H}$ let

$$V_B : S^{-1}(B_B) \rightarrow X_B \quad (28)$$

be arbitrary surjective (and measurable), with X_B arbitrarily chosen non-vacuous set such that all X_B are disjoint. Then define

$$X \triangleq \bigcup_{B \in \mathcal{H}} X_B \quad (29)$$

and define $V: \Omega \rightarrow X$ by: for any $\omega \in \Omega$, there is a unique $B(\omega) \in \mathcal{H}$ such that $\omega \in S^{-1}(B_B)$, whereupon define

$$V(\omega) \triangleq V_{B(\omega)}(\omega); \text{ all } \omega \in \Omega. \quad (30)$$

(iv) It follows from (iii) that the smallest-by subset inclusion- possible space X satisfying the required properties is of the form

$$X = \{x_B \mid B \in \mathcal{H}\}, \quad (31)$$

where each x_B represents a point which is distinct for each B .

(v) Finally, note that in (ii), $X_B = B_B$ for all $B \in \mathcal{H}$.

Proofs: Using eqs. (22), (28)-(30)

$$\hat{a} = V(S^{-1}(\bigcup_{B \in \mathcal{H}} B_B)) = \bigcup_{a \in B \in \mathcal{H}} V_B(S^{-1}(B_B)) = \bigcup_{a \in B \in \mathcal{H}} X_B \quad (32)$$

and

$$\begin{aligned} V^{-1}(\hat{a}) &= V^{-1}(\bigcup_{B \in \mathcal{H}} X_B) = \bigcup_{a \in B \in \mathcal{H}} V_B^{-1}(X_B) = \bigcup_{a \in B \in \mathcal{H}} S^{-1}(B_B) \\ &= S^{-1}(\bigcup_{B \in \mathcal{H}} B_B) = S^{-1}(C_{\{a\}}(\text{rng}(S))), \text{ all } a \in A. \end{aligned} \quad (33)$$

Eq. (24) is obtained from the injective property of Λ over A' and use of eqs. (15) and (16) in the lemma. The remainder of the results follow immediately from the constructions.

Theorem 10. Any fuzzy subset of a space has both a random subset representation-under one point coverages, and a random variable representation as the evaluation function of the random variable over some class of compound events, with the random subset and random variable naturally related.

Let B be any fuzzy subset of some given A . Thus, $\phi_B: A \rightarrow [0,1]$ is arbitrary given. Then by Theorem 1, there always exists random sets $S: \Omega \rightarrow \mathcal{P}(A)$ which are one point coverage equivalent to B . In turn, applying Theorem 9, there is a random variable $V: \Omega \rightarrow X$, with V, X , and mapping $\Lambda: A \rightarrow \mathcal{P}(X)$ all satisfying the results in Theorem 9. In particular, this implies

$$\phi_B(a) = \text{Pr}(a \in S) = \text{Pr}(V \in \hat{a}), \text{ all } a \in A. \quad (34)$$

Analogues to Theorems 2,3,4 for random variable representations of fuzzy sets have yet to be investigated.

Example illustrating random variable and random set representation of fuzzy sets.

In real world applications, classes of objects are determined by many factors where often we have no knowledge of the actual conditional and joint probabilities involved among the factors (let alone determine all of the relevant factors). In such cases, it may be both simpler and more appropriate to query experts directly as human integrators of knowledge to obtain the probabilities or possibilities of occurrences of the classes involved and their error distributions. Results from these experts show apparent probabilities for these events not adding up to one in general. In the past, normalization was carried out to make these values into "legitimate probability ones, with the implication that human error was involved in the estimates which contributed to this problem. However, since in general, the classes of interest really represent overlapping concepts, and not disjoint events, the computed values should not be expected to sum to unity. Hence, a fuzzy set membership (or possibility) function is more appropriate than a classical probability function model. Indeed, Theorems 8-10 make this idea more rigorous: the experts' responses (suitably averaged) form the fuzzy set membership function which is the same as the evaluation function of a (non-unique) random variable representing the actual elementary event space of all factors evaluated at those compound events of that space which correspond to the classes considered. The evaluation function may also be interpreted as the commonalities or one point coverage function of a (non-unique) random subset of the set of all classes- the random subset representing the possible interactions of the classes.

First, for simplicity, consider the situation described in Theorem 8: no experts used and all prob. values known:

Let $G_1^d = \{a, b, c\}$, $G_2^d = \{I, II\}$, $G_3^d = \{1, 2\}$, be the domain sets of the only factors considered for possible classes. For example, G_1 could represent lengths, G_2 , weights (heavy, light), and G_3 , shapes (of type I or type 2). Define then the classes C_1, \dots, C_6 by the following table:

Class C_j	$G_{1j} = G_1$ values for C_j	$G_{2j} = G_2$ values for C_j	$G_{3j} = G_3$ values for C_j
C_1	a, b	I	1, 2
C_2	b, c	I	1
C_3	a, c	II	1, 2
C_4	a, b	II	1, 2
C_5	a, b, c	II	1
C_6	c	I	2

Table 1.

Define the following elementary events-possible triples of values from G_1, G_2, G_3 that determine the classes:

$x_1 = (a, I, 1)$	$x_4 = (a, II, 2)$	$x_7 = (b, II, 1)$	$x_{10} = (c, I, 2)$
$x_2 = (a, I, 2)$	$x_5 = (b, I, 1)$	$x_8 = (b, II, 2)$	$x_{11} = (c, II, 1)$
$x_3 = (a, II, 1)$	$x_6 = (b, I, 2)$	$x_9 = (c, I, 1)$	$x_{12} = (c, II, 2)$

Table 2.

Thus we can tabulate for each class which elements are in it, and conversely, for each element, which classes contain it:

Class C_j	Elements in C_j	Class C_j	Elements in C_j
C_1	x_1, x_2, x_5, x_6	C_4	x_3, x_4, x_7, x_8
C_2	x_5, x_9	C_5	x_3, x_7, x_{11}
C_3	x_3, x_4, x_{11}, x_{12}	C_6	x_{10}

Table 3.

Here, $X = \{x_1, x_2, \dots, x_{12}\}$, $A = \{C_1, \dots, C_6\} \subseteq P(X)$, and the collection of all classes containing x_j is $C_{\{x_j\}}(A)$.

x_j	$C_{\{x_j\}}(A)$	x_j	$C_{\{x_j\}}(A)$	x_j	$C_{\{x_j\}}(A)$	x_j	$C_{\{x_j\}}(A)$
x_1	$\{C_1\}$	x_4	$\{C_3, C_4\}$	x_7	$\{C_4, C_5\}$	x_{10}	$\{C_6\}$
x_2	$\{C_1\}$	x_5	$\{C_1, C_2\}$	x_8	$\{C_4\}$	x_{11}	$\{C_3, C_5\}$
x_3	$\{C_3, C_4, C_5\}$	x_6	$\{C_1\}$	x_9	$\{C_2\}$	x_{12}	$\{C_3\}$

Table 4.

Next, let V_j be a random variable over G_j corresponding to the true value of factor j and \tilde{V}_j a random variable also over G_j corresponding to the observed value of factor j , for $j=1, 2, 3$. Suppose also that (V_1, V_2, V_3) are a mutually statistically independent triple, as are: $(\tilde{V}_1, \tilde{V}_2, \tilde{V}_3)$, $(V_1, \tilde{V}_2, \tilde{V}_3)$, (\tilde{V}_1, V_2, V_3) , and $(\tilde{V}_1, \tilde{V}_2, V_3)$. Suppose further that the probability functions $p(V_j | \tilde{V}_j)$ and $p(\tilde{V}_j)$, $j=1, 2, 3$, are all known, where we use the convention of identifying, where necessary, random variables and their probability functions through their typical outcomes. Define random variable $V^d(V_1, V_2, V_3)$

corresponding to the true joint factor values determining the true class, and similarly, $\tilde{V}^d(\tilde{V}_1, \tilde{V}_2, \tilde{V}_3)$ for the observed (or reported) class. Then it follows that:

$$p(V | \tilde{V} \in C_k) = \frac{\sum_{j=1}^3 \sum_{t \in G_{jk}} p(V_j | \tilde{V}_j = t) \cdot p(\tilde{V}_j = t)}{\sum_{j=1}^3 \sum_{t \in G_{jk}} p(\tilde{V}_j = t)} \quad (35)$$

$$p(C_m \text{ is true} | C_k \text{ is observed}) = p(V \in C_m | \tilde{V} \in C_k) = \frac{\sum_{j=1}^3 \sum_{t \in G_{jk}} p(V_j = t | \tilde{V}_j = t) \cdot p(\tilde{V}_j = t)}{\sum_{j=1}^3 \sum_{t \in G_{jk}} p(\tilde{V}_j = t)} \quad (36)$$

Thus, $p(C_m \text{ true} | C_k \text{ observed})$ is a computable function of C_k for each C_k in A and may be interpreted as the fuzzy set membership function or possibility function ϕ_B for possible choices of which class gave rise to the observation C_k , for any fixed k , $k=1, 2, \dots, 6$. These functions are generated by the conditional random variables $(V | \tilde{V} \in C_k)$ evaluated over the collection of compound events A . In turn, the possibility functions ϕ_B are also represented as the one point coverage functions of the random subsets $(S | \tilde{V} \in C_k)$ of A which represent the possible interactions of the classes $k=1, 2, \dots, 6$. Here, for any k , a typical outcome of $(S | \tilde{V} \in C_k)$ is some $C_{\{x_j\}}(A)$, the collection of all classes C_1, C_2, \dots

such that they interact with respect to x_j , i.e., contain x_j . Table 4 presents the ten distinct such collections of classes making up the range of $(S | \tilde{V} \in C_k)$. The probability function for this random subset is obtained from eq.(16). Some of the values (see Table 4):

$$p(S = \{C_1\} | \tilde{V} \in C_k) = \sum_{i=1, 2, 6} (p(V = x_i | \tilde{V} \in C_k)),$$

$$p(S = \{C_3, C_4, C_5\} | \tilde{V} \in C_k) = p(V = x_3 | \tilde{V} \in C_k).$$

$$p(S = \{C_3, C_4\} | \tilde{V} \in C_k) = p(V = x_4 | \tilde{V} \in C_k),$$

$$p(S = \{C_1, C_2\} | \tilde{V} \in C_k) = p(V = x_5 | \tilde{V} \in C_k),$$

$$p(S = \{C_6\} | \tilde{V} \in C_k) = p(V = x_{10} | \tilde{V} \in C_k), \text{ etc.} \quad (37)$$

which may be evaluated through eq.(35).

Now consider the converse of the above situation-where here Theorem 10 is appropriate: We know the classes of relevance $A = \{C_1, \dots, C_6\}$, but we do not know all the contributing factors and/or we do not have a handle on all the required probability functions involved in the factors; but we do have before us a panel of experts who will give the possibilities of true class values given observed ones. To simplify notation, assume that a particular C_k is chosen and all results, as above, are conditioned upon the event $\tilde{V} \in C_k$. Suppose then

We first must choose a random subset S of A (such as in eq.(37)-but here we do not know a priori the random variable V generating S and Φ). Guided by Theorem 1 and perhaps by Theorem 5, by minimal entropy or nested random set form requirements, choose $S = S_U(B)$, e.g. In this case, the probability function for S is, by defining for $m=1, \dots, 6$, (m) .

$$p(S=C^{(m)}) = \phi_p(C_-) - \phi_p(C_{-1}), \quad m=1, \dots, 6. \quad (40)$$

$$C_{i, c-1}(\text{rngS}) = \{C^{(1)}, \dots, C^{(m)}\}, \quad (41)$$

whence

$$S^{-1}(C^{(m)}) = (\phi_n(C_{-}), \phi_n(C_{-})): \quad (43)$$

(surjective) $V_{(m)}: (\phi_R(C_{m-1}), \phi_R(C_m)) \rightarrow X_{(m)}$ (44)

v defined from the v's as in eq. (30) and $v_d^d \approx v$.

SUMMARY AND CONCLUSIONS

=belj(ψor(ψ(ψsend,a(b,comm),ψvery(ψoblig(a))))))
 c#0 or so det. (os=) made over pst-24 hrs: dary & unus. long.

$$\text{tr}(c) = \sum_{(x_1, \dots, x_n \in X; n \geq 1)} (\psi_1(x_1) \cdots \psi_n(x_n)) \cdot g(D),$$

This determines the conditional $\text{tr}(d|c)$ through relation

47

1. D. Kahneman & A. Tversky, "Subjective probability: a judgment of representativeness" Cog. Psych. 3(1972), 430-454.
2. L. J. Cohen, "Can human irrationality be experimentally demonstrated?", Behav. & Brain Sci. 4(1981), 317-370.
3. T. L. Fine, Theories of Probability, Acad. Press, N.Y., 1973
4. A. N. Freeling, Alternative Theories of Belief and Implications for Incoherence, Reconciliation, and Sensitivity Analysis, Decision Science Consortium, Inc., Tech. Report No. 81-4, Nov., 1981.
5. E. G. Manes, "A class of fuzzy theories", J. Math. Anal. & Applic. 85 (1982), 409-451.
6. I. R. Goodman & H. T. Nguyen, "Uncertainty modeling and a possibilistic approach to parameter estimation", submitted to J. Math. Anal. & Applic., 1983.
7. I. R. Goodman, "Fuzzy sets as equivalence classes of random sets", in Recent Developments in Fuzzy Sets and Possibility Theory, R. R. Yager, ed., Pergamon Press, New York, 1981, 327-342.
8. G. Shafer, "Lindley's paradox", J. Amer. Stat. Assoc. 77, No. 378 (1982), 325-351.
9. R. V. Brown & D. V. Lindley, "Improving judgment by reconciling incoherence", Theory & Decis. 14(1982)113-132.
10. I. R. Goodman, "Some fuzzy set operations which induce homomorphic random set operations", Proc. 26th Conf. Gen. Sys. Res. (AAAS) 1982, 417-426.
11. H. T. Nguyen, "On random sets and belief functions", J. Math. Anal. & Applic. 65(1978), 531-542.
12. U. Hühle, "Representation theorems for L-fuzzy quantities", Fuzzy Sets & Sys. 5 (1981), 83-107.
13. I. R. Goodman, "Identification of fuzzy sets with random sets", survey paper, accepted for Encyclopedia of Systems & Control, M. Singh, ed., Pergamon Press, New York (to appear, 1984).
14. E. P. Klement, "Operations on fuzzy sets and fuzzy numbers related to triangular norms", Proc. 11th Inter. Symp. Mult. Logic. (1981), 218-225.
15. I. R. Goodman, "Characterizations of n-ary fuzzy set operations which induce homomorphic random set operations", in Fuzzy Information and Decision Processes, M. M. Gupta & E. Sanchez, eds., North-Holland, New York, 1982, 203-212.
16. I. R. Goodman, "Some new results concerning random sets and fuzzy sets", submitted to Info. Sci., 1983.
17. I. R. Goodman, "Some asymptotic properties of fuzzy set systems", Proc. 2nd World Conf. Math. Serv. Man, Can. Is., 1982, 312-317.
18. H. Dishkant, "About membership function estimation", Fuzzy Sets & Sys. 5 (1981), 141-147.
19. I. R. Goodman, "An approach to the data association problem through possibility theory", Proc. 5th MIT/ONR C³ Sys., Dec., 1982, 209-215.
20. I. R. Goodman, "PACT: possibilistic approach to correlation and tracking", Proc. 16th Asilomar Conf. Circ. Sys. Comp. (IEEE), Nov., 1982, 359-363.
21. I. R. Goodman, "Applications of fuzzy set theory to parameter estimation and tracking", Proc. Amer. Contr. Conf. 1983, 1236-1237.
22. A. M. Pitts, "Fuzzy sets do not form a topos", Fuzzy Sets & Sys. 8 (1982), 101-104.
23. J. C. Carrega, "The categories Set H and Fuz H", Fuzzy Sets & Sys. 9 (1983), 327-332.
24. D. Ponsse, "Some remarks on the category Fuz H of M. Eytan", Fuzzy Sets & Sys. 9 (1983), 199-204. For Eytan's controversial article, see Fuzzy Sets & Sys. 5 (1981), 47-67.
25. M. P. Fourmann & D. S. Scott, "Sheaves and logic", Lecture Notes in Math. No. 753, Springer-Verlag, New York, 1979, 302-401.
26. J. A. Goguen, "Concept representation in natural and artificial languages: axioms, extensions, and applications for fuzzy sets", Int. J. Man-Machine Stud. 6 (1974), 513-561.
27. H. T. Nguyen, I. R. Goodman & W. Ruitenberg, "Additional relations between Fuz(H) and Higg(H)", to be pub. 1984.
28. G. Matheron, Random Sets and Integral Geometry, John Wiley & Sons, New York, 1975.

SOFTWARE FOR EXPLICITLY PROBABILISTIC MATHEMATICS

Steven N. Goldstein

The MITRE Corporation
1820 Dolley Madison Boulevard
McLean, Virginia 22102

Introduction

The conventional way to estimate quantities for decision making purposes is to employ informed judgment about the applicable estimating relationships (equations, algorithms, models) and the values of their constituent parameters. More often than not, one or more key parameters are not known with certainty, and informed judgment takes the form of a subjective point estimate ("guesstimate"). Many practitioners are content to use the resulting point estimate for the result. The more diligent will re-do the calculation combining estimates of upper and lower bounds of the input parameters to estimate upper and lower limits for the result. Such diligence is facilitated by the increasingly convenient spread sheet computational packages for personal computers.

This paper describes software for incorporating uncertainty directly and explicitly into the calculations to yield a probability distribution as the result rather than a set of point estimates. The software permits the user to represent the uncertain parameters by any of about sixteen probability density functions. This is done with simple assignment statements like:

'A' IS 'NOR 10 1'

which represents the parameter A as being normally distributed (bell-shaped curve) with a mean of 10 and a standard deviation of 1. The parameters represented by probability distributions rather than point estimates are called probabilistic variables. The software combines the probabilistic variables in expressions almost like those involving deterministic variables, except that instead of using the familiar arithmetic operators (+, -, x, /, *) corresponding primitive functions like PL ("Plus", for addition), MI ("Minus", for subtraction), TI ("Times", for multiplication), DI ("Divided by", for division), and EX for exponentiation are substituted.¹ Thus, the expression

$$D = (A + B)C,$$

is written as

$$D \leftarrow (A \text{ PL } B) \text{ EX } C,$$

assuming that A, B and C have all been assigned definitions as probabilistic variables. And the result, D, is itself a probability distribution rather than a point estimate.

A probability distribution provides much more information than a single estimate, or even an estimate with upper and lower bounds. For example, the P-th centile, the value which is greater than P-percent of the other possible values of the result, can be found directly from the probability distribution of the result.² Similarly, summary statistics or graphical representations of the probabilistic result can also be produced.³

Perhaps one of the most useful features is the ability to compare probabilistic results which represent competing alternatives. Often, the expected result (mean of the probability distribution) is independent of the amount of variability (spread, dispersion) of the result. It is not unusual to find circumstances where the alternative that has the greatest expected value also has the highest dispersion. Greater dispersion can mean greater risk; in fact risk is often described as being directly proportional to some (exponential) power of a measure of the dispersion such as the variance -- the bigger the spread, the greater the risk, by some definitions.⁴ The decision maker may place a premium on risk reduction and may be willing to trade off expected performance in order to reduce the risk of the system's not performing adequately in adverse situations.⁵

So, to take stock before moving on, this paper will describe a computerized method for representing uncertainty explicitly in calculations concerning individual alternatives and for making choices among several uncertain alternatives.

Related Work

I credit the early motivation for the work to my research with M. Granger Morgan at Carnegie-Mellon University.⁶ Morgan has made many practical contributions to the art of dealing explicitly with uncertainty in public policy issues with high technology content. He and his colleague, Max Henrion, have developed a software system with similar capabilities which they call DEMOS. They are readying DEMOS for general use with support from the National Science Foundation. DEMOS is coded in Pascal, a higher order computer language which is viewed by many as a root of the new Ada language.

In other related work, Leo H. Groner of IBM⁷ has been investigating a number of computerized methods for performing explicitly probabilistic calculations. Some are similar to the approach I have taken, though Groner uses a different representation of probability density functions. Much of his work is also coded in APL.

Also, Bonner and Moore Associates Inc., Houston, Texas⁸ offer a commercial FORTRAN-based software package for probabilistic calculations. The calculations are performed with Monte-Carlo simulation techniques. They have named it PAUS. PAUS is intended for a variety of business applications. It also appears adaptable to engineering applications.

My own system is in an advanced development stage, ready for application testing. It is based on a numerical approach which is similar to the

convolution of probability distributions. My objective is to enable analysts to systematically and explicitly incorporate uncertainty into their calculations with a minimum of set-up overhead.

An Illustration of an Explicitly Probabilistic Calculation

A sample problem solution will be presented to illustrate both the way the software can be used and the extra dimension of information that is available from following the propagation of uncertainty through a calculation. Consider a system with two types of elements with unit operating costs A and B, respectively. The total system operating cost is described by the formula

$$D = (M \times A + N \times B)C,$$

where M is the number of units of the first type in the system, and N is the number of units of the second type. C is an exponent which represents the influence of scale on costs, i.e., economy (C less than 1) or diseconomy (C greater than 1) of scale.

Imagine one system composed of one unit of the first kind and 35 units of the second kind. Its operating cost would then be:

$$E = (A + 35 \times B)C.$$

A second system alternative with no units of the first kind and 75 of the second kind would have the following operating cost:

$$F = (75 \times B)C.$$

The first type of item could very well be a large central component such as a mainframe host computer and the second item type could be a smaller remote counterpart like a local processor. Actual cost equations would be more complicated, but similar in principle to the illustration.

Suppose an analyst had made (or had obtained) the following point estimates: A = 10; B = 0.24; C = 0.63.

The cost estimates would then have been:

$$E = (10 + 35 \times 0.24)^{0.63} = 6.26, \text{ and}$$

$$F = (75 \times 0.24)^{0.63} = 6.18.$$

These differ by 1.3 percent, certainly too close a call for a clear preference.

Suppose, further, that instead of just point estimates, the analyst was able to encode the uncertainty in the estimates of the parameters with the following probability density functions (PDFs):

A ~ NORMAL (mean = 10; standard deviation = 1)

B ~ GAMMA ("shape" = 2; "scale" = 7)

C ~ TRIANGULAR (low = 0.575; apex = 0.6; high = 0.75).

The probabilistic parameters will yield a distribution rather than a point estimate for operating costs E and F. This is how it is done with the software:

1. Enter the PDFs for A, B and C, as shown in Exhibit 1.
2. Enter the expressions for E and F, as shown in Exhibit 2. Where probabilistic variables are combined, substitute the functions "PL" for addition, "MI" for subtraction, "DI" for division, "TI" for multiplication, and "EX" for exponentiation. Note that in an operation involving a constant and a

probabilistic variable, the regular mathematical operator may be used (e.g., 35 x B).

3. The statistics of the resulting distributions are generated with "DSTAT" and illustrated with "BOXPLOT". Also, vertical boxplots for the two results, E and F can be generated side-by-side (to the same scale) with the function "COMPARE" [McNeil, 1977] as shown in Exhibit 2.

Referring to Exhibit 2, one sees that although the two alternatives have about the same expected operating cost (as was predicted by the original point estimates and borne out by the medians and means of the probabilistic results) alternative F has about twice the dispersion (as measured by the standard deviation) as alternative E. So, the choice between the two depends on whether the objective is to minimize uncertainty (E is better), minimize the highest possible cost (E is better), minimize the lowest possible cost (F is better), or some other objective. The software includes a "DECISION" function which automatically selects the best alternative. Its use is illustrated in Exhibit 3, in which the operating costs E and F have first been subtracted from an arbitrary revenue of 12 to yield an operating profit. This is because the decision rules are designed for maximization rather than minimization.

Example 1: Probabilistic Inputs

A:

A=127 MEAN 10 1

BOXPLOT A

7.8

12.4

DSTAT A

MINIMUM: 7.7687 MAXIMUM: 12.3916 RANGE: 4.60294
SAMPLE SIZE: 127 MEDIAN: 9.89087
MEAN: 9.94749 VARIANCE: 1.07358 STD DEVI: 1.03614
MEAN DEVI: 0.875871 SKEWNESS: -0.0252866 KURTOSIS: 2.18652

B:

B=127 GAM 2 7

BOXPLOT B

0

1.3

DSTAT B

MINIMUM: 0.043226 MAXIMUM: 1.36275 RANGE: 1.25953
SAMPLE SIZE: 127 MEDIAN: 0.242528
MEAN: 0.28299 VARIANCE: 0.0400479 STD DEVI: 0.20012
MEAN DEVI: 0.143766 SKEWNESS: 1.79176 KURTOSIS: 7.92908

C:

C=127 TRI .575 .6 .75

BOXPLOT C

0.6

0.7

DSTAT C

MINIMUM: 0.576573 MAXIMUM: 0.741978 RANGE: 0.165403
SAMPLE SIZE: 127 MEDIAN: 0.631359
MEAN: 0.641004 VARIANCE: 0.00152142 STD DEVI: 0.0390054
MEAN DEVI: 0.0321159 SKEWNESS: 0.583786 KURTOSIS: 2.42062

Exhibit 1

Example 1: Summary of Results

E: R (A PL 35KB) EX C

BOXPLOT E

4.3 12.6

MINIMUM: 4.2501 MAXIMUM: 12.5706 RANGE: 8.32053
 SAMPLE SIZE: 127 MEDIAN: 6.40701
 MEAN: 6.71906 VARIANCE: 2.51455 STD DEV: 1.58573
 MEAN DEV: 1.2236 SKEWNESS: 1.07411 KURTOSIS: 4.25071

F: P (75KB) EX C

BOXPLOT F

2.1 17.7

MINIMUM: 2.12909 MAXIMUM: 17.7175 RANGE: 15.5884
 SAMPLE SIZE: 127 MEDIAN: 6.31663
 MEAN: 6.74461 VARIANCE: 9.24134 STD DEV: 3.03996
 MEAN DEV: 2.32959 SKEWNESS: 0.950114 KURTOSIS: 4.0354

ME: [1.11]
 COMPARE M

17.7175
 2.12909

← Medians

E F

Example 1: Apply Decision Rules to E and F Operating Profits				
DECISION @ (12-M)				
DECISION STRATEGIES *				
MAXIMUM	:	CHOOSE SYSTEM(S) 2 WITH VALUE=9.87091		
MINIMUM	:	1	:	=0.570625
MINIMUM REGRET	:	1	:	=2.21406
MIN EXPECTED REGRET:	:	1	:	=0.539302
MAX EXPECTED VALUE :	:	1	:	=5.28094
MIN VARIANCE	:	1	:	=2.51455
* Systems 1 = E; 2 = F				
Exhibit 3				

50

A Command and Control Application: Message Processing Performance

Robustness in command and control (C²) networks is often enhanced by redundancy of message pathways. The links of any path may encompass a diversity of frequencies, transmission media, and formats (e.g. voice, telemetry). A variety of processes occur at nodes to convert messages for retransmission where content, format, transmission medium, frequency, or other significant properties of the message may change. Exhibit 4 depicts redundant paths through a hypothetical message processing network. The objective is to transmit an authorized message (abbreviated "MSG" in the exhibit) to a recipient who can determine that it is valid and who can take the actions(s) indicated by the message.

The paths, or subsystems, are identified as follows:

- o On Path A, the message is processed for transmission on a data network (Process P1), is transmitted (P3), received and validated (P10);
- o on Path B, it is processed (P2), transmitted (P4), received and recoded for voice

All of the processes are assumed to be statistically independent of each other. Nominal performance statistics, which are also hypothetical though not unrealistic, are given in Exhibit 5. I chose the Gamma Distribution to represent the performance statistics, because the Gamma Distribution is unimodal and positively skewed, is defined over the domain zero to infinity and has a mean and standard deviation which are simple functions of the parameters of the probability density function.¹⁰ The parameters of the distributions for each independent process are given in the right-most column in Exhibit 5. Summary graphical comparisons of the individual processing times are shown in Exhibit 6.

Example 2: Network Message Processing

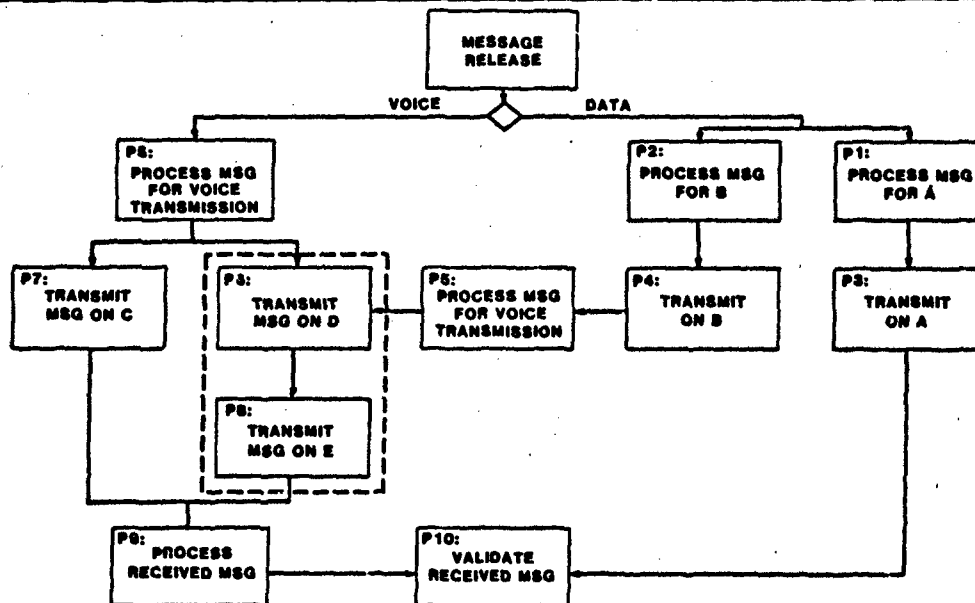


Exhibit 4

transmission (P5), transmitted as voice (P6) with real-time relay through E, received and processed (P9) at the destination and validated (P10);

- o on Path C, the message is processed for voice transmission (P5) and sent (P7) directly to the destination where it is received, processed (P9), and validated (P10);
- o finally, on path D it is processed for voice transmission (P5) as in path C, but must pass through and intermediate receipt and transmission stage involving back-to-back processing (P6 followed by P8) before being received, processed (P9) and validated (P10) at the destination.

The end-to-end transit times for messages are the sums of the processes along each path: $(P1 + P3 + P10)$, for Path A; $(P2 + P4 + P5 + P6 + P9 + P11)$, for Path B; $(P5 + P7 + P9 + P10)$, for Path C, and $(P5 + P6 + P8 + P9 + P10)$, for Path D. Because each process is treated as a probabilistic variable, each path sum is found using the probabilistic summing function, PL, instead of the plus sign. The assignment statement which sets up the sum for path D is shown in Exhibit 7 along with graphical and statistical analyses of the results. Note the smooth envelope of the histogram-like stemleaf plot of the end-to-end time. The plots for the other paths are similar. Summary graphical comparisons of the end-to-end transit times are shown in Exhibit 8.

Example 2: Nominal Processing Time Statistics

ACTIVITY	MEAN, STANDARD DEVIATION (min)	PARAMETERS OF BEST FIT GAMMA * DISTRIBUTION E, L
P1: Process Message	6, 3	4.0, 0.67
Process Message For B	7, 2	12.75, 1.75
P5: Process Message For Voice Transmission	6, 4	2.25, 0.375
P3: Transmit Message On A	3, 1	9.0, 3.0
P4: Transmit Message On B	3, 2	2.25, 0.75
P7: Transmit Message On C	5, 4	1.50, 0.31
P6: Transmit Message On D	6, 5	1.44, 0.24
P8: Transmit Message On E	4, 5	1.0, 0.25
P9: Process Received Message	6, 2	9.0, 1.5
P10: Validate Received Message	1, 1	1.0, 1.0

$$f(x) = \begin{cases} \frac{L}{\Gamma(E)} x^{E-1} e^{-Lx} & , x \geq 0 \\ 0 & , \text{elsewhere} \end{cases}$$

Exhibit 5

Example 2: Summary of P1-P10 Inputs




Exhibit 6

Example 2: End-To-End Tiring For Path D




Exhibit 7

Example 2: Summary of End-To-End Paths A-D Transit Times




Exhibit 8

Here are two examples of the usefulness of the output formats. First, note that the 90th centile for path D is 32 minutes. Similarly, the 90th centiles for the others are: 14 minutes for A; 37 minutes for B, and 26 minutes for C. These results are read from the tabulations immediately beneath the cumulative distribution graphs (CPlots) as in Exhibit 7. Second, the output graphics can be used to estimate overall system performance. For example, Bullers has offered the following analysis:¹¹ the probability that the message gets through Path D within 15 minutes is 0.14 (graphically, from Exhibit 7). Therefore the probability that it does not get through within 15 minutes is $1 - 0.14 = 0.86$. Likewise, the probabilities for not getting through Paths A, B and C within 15 minutes are 0.08, 1, and 0.69 respectively. Therefore, the probability that no message gets through any path within 15 minutes (assuming independence among processes) is $0.08 \times 1 \times 0.69 \times 0.86 = 0.05$. Thus, the probability that a message does get through the system within 15 minutes is $1 - 0.05$, or 0.95.

These types of output could also prove useful in assessing impacts of upgrading processes in the network. With minor extensions that normally require discrete event simulation, several other interesting avenues of inquiry could be pursued; time-dependent survivability, for example. However, with one small exception, the probabilistic variables must be assumed to be independent of each other if the approach taken by this software is to apply. Selected characteristics of the software, including the assumption of independence, are discussed in the concluding section.

Description of the Software

How It Works

The probability density function (PDF) of a variable which is the sum of two probabilistic variables is the mathematical convolution of their PDFs. Taking a convolution is like scanning one PDF across the other, exposing every point in the domain of one probabilistic variable to every point in the domain of the other one.

This software works in a similar way: each probabilistic variable is represented by a vector of N numbers which constitutes a sample from the universe of numbers with the indicated PDF. When two variables are combined by addition, subtraction, multiplication, etc., all possible combinations (additions, subtractions, etc.) of the elements of the vectors representing the two variables are formed. This yields a table of dimension $N \times N$, Exhibit 9. Because the table is too large to carry around in the computer for further calculations, an N -element sample is taken from the table. The sample is formed by sorting the N^2 elements of the table, finding the median and bisecting the sorted set at the median. Then the lower and upper halves are bisected at their respective medians; then each quartile is bisected, and so forth, until a sufficient number of medians result. The number of fractiles equals 2^P , where P is the number of successive bisections. Therefore, the number of medians (the boundaries between adjacent fractiles) is $2^P - 1$.

Once two variables are combined through $N \times N$ expansion followed by median bisection shrinking, the result is available for combining in like manner with another variable (for example, in the sum $A + B + C$, B and C are combined first -- APL evaluates expressions from right to left -- and the result is then combined with A). It is convenient to keep the vectors the same length throughout this process. For $P = 7$, there are 127 elements in the vector. If N is chosen to be 127, N^2 is 16,129 entries. A table this size uses a lot of main memory -- over 120 kilobytes for 8-byte floating point numbers. (For $P = 8$, N^2 is over 65,536 requiring over 520 kilobytes of storage.) Thus, for practical purposes, the sample size for each probabilistic variable is 127.¹²

Why bisect by medians? In part, it is a philosophical throwback to methods for eliciting subjective probability distributions for uncertain parameters. The subject is asked to give an estimate such that the actual, but unknown, value is equiprobably above or below the estimate. This bisects the probability distribution at the median. Then the subject is told to assume that the actual value lies in the lower half of the bisected distribution and is asked to estimate a value such that the actual value is equiprobably above or below it. This estimate bisects the lower half. The same operation is done with the upper half, yielding the boundary values of the quartiles of the subjective probability distribution. The process can be repeated to yield octiles and so forth (at the risk of exhausting the subject). Because many of the inputs to probabilistic computations with this model will be subjective estimates of uncertainty rather than best-fits to empirical data, there is a consistency in using the median bisection technique to shrink the table down to manageable proportions. Also, because many of the operations are inherently non-linear, it makes more sense to bisect by medians than by alternatives such as means (should one use geometric means for operations involving multiplications, for example?).

However, by selecting only the medians, the extreme tails of the N^2 -member distribution are left out, e.g. there will be 125 entries in the 16,129-element table that will be less (greater) than the lowest (highest) number in the 127-element sample of medians. This omission may be undesirable to analysts who focus on the extreme values of a distribution. When the results of using this software were compared with theoretical calculations, the standard deviations often underestimated the theoretical values by about 5 or 6 percent. This could be related to truncating the tails during median bisection. The discrepancy might be mitigated by including the two end points to yield a $2^P + 1$ or 129-element sample, but I have not yet investigated this possibility.

What Capabilities it Includes

This is a general purpose probabilistic mathematics system. The functions which should prove useful to general users are listed in Exhibit 10. They are described below.

Continuous PDFs. These represent continuous random variables by sets of numbers sampled from the universe of appropriately distributed numbers. The density of sampled numbers between any two points in the domain of the PDF is roughly proportional (within sampling accuracy) to the area under the PDF between those two points. Sample generation starts with pseudo-random numbers. These are then transformed into the appropriate distributions by a variety of algorithms which include inverse transformations and acceptance-rejection methods.¹³ As was illustrated in the examples, it is helpful to display the sample with some combination of a boxplot, stemleaf and statistical summary to make sure that it fits the actual PDF accurately (for example, is the sample mean within a few standard errors of the theoretical mean?). If not, new samples should be generated until one that satisfies the user's needs is obtained. Checks like these could be built into a production version of the software, but for now they must be done manually.

Discrete PDFs. These represent discrete random variables (e.g., integers) in a way similar to the continuous PDFs just described.

Primitives. These do probabilistic addition, subtraction, multiplication, division and exponentiation. They work by computing a numerical analog of a convolution, as described above.

Display Functions. These produce statistical summaries and graphical representations of one or more samples. Their use has been illustrated in the examples. Three of these functions merit special mention: "STEMLEAF", "BOXPLOT" and "COMPARE", all of which were taken from [McNeil, 1977].

The stemleaf plot appears visually similar to a histogram with a vertical axis, but it contains more information about the numbers in the sample. "STEMLEAF" divides the sample range into intervals and prints two leading significant digits of each interval to the left side of the vertical axis. A digit equal to the (rounded) third significant figure of each number in the sample is printed to the right of the axis and aligned with the two-digit number representing the proper interval. Thus, in addition to giving a visual impression of how the sample is distributed, the stem leaf plot shows the actual (rounded) numbers.

Software Description: How It Works




Exhibit 9

The "BOXPLOT" function produces an alternative visual impression of the sample distribution. Imagine circling an invisible mark on the horizontal or vertical axis for each member of the sample. Then enclose the range from the 25th to the 75th percentiles (the interquartile distance) in a box. Denote the median (50th percentile) by an asterisk (horizontal boxplot) or a hyphen (vertical boxplot) placed on the axis inside the box. Next, place two x's on the axis to mark the sample members furthest from, but within one interquartile distance of either side of the box. Note with open circles each sample member lying outside the x's (denote multiple values, where two or more circles overlap within the resolution of the inter, with an adjacent numeral). Identify as outliers those circles lying more than 1.5 interquartile distances outside the box by filling in the circles. Significantly skewed distributions are easily detected by the lack of symmetry (of the median, x's and circles) with respect to the box. Highly peaked distributions will have few circles; diffuse distributions should have several. This description for a boxplot was adapted by McNeil from John Tukey's original version.

"COMPARE" prints side-by-side vertical boxplots, all to the same scale. The upper and lower bounds of the total range of the samples are printed at the top left. "COMPARE" facilitates a visual comparison of several samples (recall that 10 processing time distributions were compared in Exhibit 6) by depicting the relative positions of their medians, the dimensions of their dispersions, and their ranges.

Correlation. By taking all possible combinations of variables in samples in $N \times N$ expansion, the primitive algorithms destroy any correlation between probabilistic variables. Sometimes it is desirable to maintain a dependency among variables in an equation. "CORR" is a function which is patterned after Bonner and Moore Associates' correlation algorithm. It allows the user to generate a normally

distributed sample (or to approximate other symmetric distributions) which is correlated with one or two already generated samples (which may, but need not, be correlated with each other). The user specifies the other probabilistic variables and the desired correlation coefficients as arguments of the function. The correlation coefficients must satisfy certain constraints, and "CORR" checks for this.

Once the correlated sets have been generated, the order of their elements must be preserved, and the $N \times N$ table expansion cannot be used. They can be combined in APL just by using the native primitive operators (+, -, x, +, and * for exponentiation) which combine the N elements of each sample on a one-on-one basis. The resulting N-element set will incorporate the influence of the dependency among its variables. It can then be combined freely with other non-correlated variables using the regular functions already described.

The function "RHO" calculates the correlation coefficient between any two N-element samples. "MAT" displays the moment matrix, covariance matrix, and correlation matrix for N-element samples. "RHO" and "MAT" can be used to test for or to verify correlations among variables.

Miscellaneous Functions. The six functions shown are machine-dependent. They are coded for IBM's VS APL running under the Conversational Monitor System (CMS). With implementation-dependent modifications, however, they can be adapted to versions of APL running on other systems. The most important function of this group is "IS". It takes the place of APL's left-pointing assignment arrow, and it permits use of this software system from a non-APL keyboard. "RESET" also performs an operation that would require a special APL character, the right-pointing arrow. It clears out any calculations which may have been suspended as well as functions which may have been pendant (awaiting intermediate results) when an error

is encountered and calculation stops. And it also saves the cleaned-up workspace.

There are several other included functions which are called by the functions listed in Exhibit 10, but which are of no direct use to the user. For example the placement of the graphic characters in the displays generated by "BOXPLOT" and "COMPARE" are computed by the function "FILL".

SOFTWARE SUMMARY

Probability Density Functions (PDFs)

Continuous PDFs

BET Beta
CHI Chi-squared
EXP Exponential
GAM Gamma
LOG Log normal
NOR Normal
RAY Rayleigh
(circular normal)
TRA Trapezoidal
TRI Triangular
UNI Uniform
WEI Weibull

Discrete PDFs

BIN Binomial
GEO Geometric
HYP Hypergeometric
NBI Negative Binomial
POI Poisson

Primitive Functions

DI Division ("divided by")
EX Exponentiation
MI Subtraction ("minus")
PL Addition ("plus")
TI Multiplication ("times")

Display Functions

BOXPLOT Graphic representation of a sample distribution
CENT Centiles of a sample of numbers
COMPARE Side-by-side comparison of several boxplots
CPLOT Plot of cumulative distribution of a sample
DECISION Applies decision rules and identifies best alternatives
DSTAT Summary statistics of a sample
FREQ Unique members of a sample and their frequency of occurrence
PLOT An X-Y plot, scattergram
STEMLEAF A histogram-like representation of a sample

Other Functions

Correlation

CORR Forms normally distributed sample with prescribed correlation to one or two other samples
MAT Moment, covariance and correlation matrices among samples
RHO Correlation coefficient between two samples

Miscellaneous (Implementation-Dependent) Functions

DTG Prints date and time
IS Implements assignment operator for non-APL keyboards
RESET Clears out suspended operations and saves the workspace
SKIP Creates blank lines on the output display
STATUS Shows CPU use and workspace status
TIME Times the execution of a function

Exhibit 10

Concluding Remarks

The motivation for creating this software grew from my interest in examining and quantitatively describing the influence of uncertainty on making decisions -- especially when an explicit treatment of uncertainty and an evaluation of the resulting risk reverses the decision that would otherwise have been made. My intent was to develop a general purpose system that requires little set-up time to run a problem and in which the problem can be entered in a form that is natural to analysts.

The system grew by pieces as a spare-time interest over the past 18 months. Now, it is in a development state and ready for testing and refinement with actual applications. I would welcome the opportunity to collaborate with investigators who seek a way of systematically evaluating the role of uncertainty in their quantitative analyses.

The graphics are of a homespun variety. One tends to appreciate their richness with use, and to forgive the unpredictability of their embedded scaling algorithms. Yet, with all the commercial graphics software on the market, the possibility of adding some gloss and sparkle to the display capabilities is enticing. Suggestions and opportunities for collaboration on the display features are therefore also invited.

Footnotes

1 This particular implementation of the software has been coded in the APL programming language because APL has built-in array processing features that facilitate the necessary computations. In APL, the assignment symbol is the back-arrow \leftarrow ; multiplication, is \times ; division, is \div , and exponentiation is a single asterisk, $*$. However, beyond knowing these few symbols and the precedence relationships for combining operators in APL, the user need not have a working knowledge of APL.

2 The software includes a function "CENT" to extract this information. Its syntax is P CENT X, where P is a list containing one or more percents of interest. The output contains the corresponding centiles of the probabilistic result X.

3 These capabilities are also included in the software: a function which compiles summary statistics and several functions which produce graphical representations of results on the printer are included.

4 Another closely related definition of risk considers only that portion of the dispersion that falls to one (undesired) side of a critical threshold such as profit-loss breakeven point, yield strength of material, toxicity level of a dangerous substance, etc.

5 Decision rules which balance expected outcomes with risk are called mean-risk decision functions. The software includes a function which applies a number of decision rules (including mean-risk and game-theoretic rules) to probabilistic alternatives and indicates the best choice for each rule. It also includes graphical comparison capabilities.

6 Department of Engineering and Public Policy, Carnegie-Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, (412)578-2672.

7 Leo H. Groner, IBM, Inc., Box 390, Poughkeepsie, N.Y. 12602, (914)463-3615.

8 Alan Jenkin, Bonner and Moore Associates, Inc.,
2727 Allen Parkway, Houston, TX 77019;
(713)522-6800.

9 Notes: (1) The left-pointing arrow is used for assignment in the APL language; the function "IS" has also been provided for use on keyboards which do not have the arrow; (2) The number 127 instructs the computer to represent the probabilistic variable by a vector of 127 variates -- in a later, more refined version of the software, it should no longer be necessary to specify the number of variates; (3) The boxplot [McNeil, 1977] which is a representation of the distribution--the median is represented by the asterisk, and outliers are open (near) or filled-in (far) circles -- and the statistical description invoked with the function "DSTAT" [Ramsey, 1981] are not obligatory; they are available to characterize and check the statistical samples which have been generated to represent the probabilistic variables.

10 If the parameters are E and L, the mean is E/L , and the variance is E/L^2 . Also, the mode is given by a simple expression, $(E-1)/L$, which adds to the convenience of using the Gamma Distribution.

11 Bullers, J. W., "Methodology for Calculating the Expected Performance of European Theater Communications Systems", working draft dated 31 March 1983, The MITRE Corporation, McLean, VA 22102

12 Even this manageable number requires a 1 Megabyte virtual machine on an IBM 370 mainframe running under CMS.

13 See, for example, [Fishman, 1978].

14 For example, one can type 'X' IS '127 NOR 5 3' (enclosing the left and right arguments in single quotes). This is equivalent to the assignment statement $X \leftarrow 127 \text{ NOR } 5 \text{ } 3$.

Cited References

Fishman, George S., "Principles of Discrete Event Simulation," John Wiley, New York, 1978.

McNeil, D. R., "Interactive Data Analysis," Wiley Interscience, New York, 1977.

Ramsey, J. B. and Musgrave, G. L., "APL-STAT, A Do It Yourself Guide to Computational Statistics Using APL," Lifetime Learning Publications, Belmont, CA 1981.

SECTION II

Humans: Models and Interfaces

Preceding page blank

COMMAND DECISION MAKERS AND THEIR MODES OF INTERACTION

P. D. Morgan

Scicon Ltd., 49 Berners St., London W1P 4AQ

Abstract. An analytic model of the basic decision making process is presented; in this it is shown how external inputs may be employed to reduce the scale of the decision making task, which is represented in the form of a pair of coupled decision making processes. An approach based on control and estimation theory is used in modelling the decision making process; this is supplemented by concepts drawn from knowledge representation and from game tree search. Finally the decision maker is examined as a part of a control process in order to identify the relative tempos of different operations and interactions.

1. INTRODUCTION

The Commander and The Command System are subjected to a steadily increasing load due to the continuing increase in the range of the combat horizon and the steady development of Counter C² techniques. The solution generally adopted to support this load has been to resolve the Commander's objectives into ranges of more limited objectives, and to assign these to subordinate commanders; however this introduces problems of co-ordination between these subordinates [1], which at present are handled entirely manually. The introduction of automation into the Command System has made apparent the need for a quantitative understanding of the command decision process and of the interactions between decision makers; this understanding is required to provide a basis for the analysis and design of Command Systems.

In previous work the operation of the Command System was represented in terms of a basic module; this included the decision making process together with its means of interacting with its environment [2]. In this approach the problems of controllability and observability [3] were addressed by means of the nesting of modules to reflect the hierarchical structure, and the use of 'Volumes of Interest' and their overlaps to determine lateral communications [1], [4].

A further conclusion of this work was that decision makers within the Command System can be classified in terms of goal-seeking behaviour [5], [6]; the three classes used correspond to: goal seeking behaviour; multi-goal seeking, adaptive behaviour; and ideal seeking, learning behaviour. On this basis it is apparent that the decision maker must be provided with

- i) resources usable to explore and modify the environment;
- ii) finite sets of population and control models representing the contents of the external environment and the operation of his resources within the environment;
- iii) a goal or goals which may be represented in

terms of some objective function, or of some goal state;

- iv) a search strategy for evaluating and selecting options in terms of the current and goal states.

Examination of these requirements in terms of control and estimation theory, and of game tree search [7] shows that a decision maker with finite data processing capacity requires some method of truncating the estimation and search processes when handling other than trivial tasks. These truncations carry with them the risk of excluding appropriate decisions; consequently it is necessary that the decision maker be capable of adapting the truncation processes in accordance with past experience and the current state of the environment.

The model of the adaptive decision maker is based on the two stage model of Boettcher and Levis [8] and on the model previously discussed [2]. The basic decision making process consists of the determination of the state of the environment (Situation Evaluation) and the choice of an appropriate response (Response Allocation). Adaptive behaviour is obtained by modelling the decision maker in terms of two parallel decision processes; an 'on-line' decision process concerned with assessing the environment and directing the operation of the resources, and an 'off-line' decision process concerned with optimising the truncation of the estimation and search processes.

This model was developed to explore the various modes of interaction within a Command System, and as an aid to determining appropriate man/machine roles within a Command System. An approach based on control and estimation theory was used as it facilitated examination of the internal and external interactions of the decision maker.

2. THE BASIC DECISION MAKING PROCESS

The decision maker is envisaged as operating in an environment in which a variety of distinguishable phenomena may occur, either serially or in parallel. The phenomena are represented to the decision maker in the form of a time-ordered sequence of noisy measurements, x , which are evaluated to determine their individual significances by the Situation Evaluation element; the output from Situation Evaluation takes the form of a sequence of estimates, z , of elements of the situation related to individual phenomena, or to associated groups of phenomena. In the Response Allocation element this sequence of estimates is used to determine the discrepancy between the goal state, Z_G , and the current state, Z_0 , and to select appropriate control outputs, y ; these outputs correspond to the allocation of particular responses to the various resources controlled by the decision maker.

In this case it is assumed that each phenomenon is a member of one of a finite but possibly large set of




Fig. 1 Situation Evaluation

populations or classes, and that the decision maker has access to a finite but possibly large set of responses. Each population has an associated population model for the generation of the individual situation estimates, z , on the basis of the measures, x ; similarly each response has associated control models for the generation of the control outputs. Thus the Situation Evaluation element is concerned with the identification of the population from which the measure, x , is drawn, followed by the application of the relevant population model; and the Response Allocation element is concerned with the evaluation and selection of a particular response, followed by the application of the associated control model.

The tasks of searching through these range of these populations and responses will result in an unacceptable processing load; consequently some means of truncating the searches is required. The approach adopted is similar to the 'frame' approach to knowledge representation [7]; it assumed that in any given period the populations observable will be members of a predefined subset of the set of all known populations; in addition, the available responses will form a subset of the set of possible responses; in combination this pair of subsets corresponds to a frame. The frame to be employed is defined for the decision maker on the basis of the expected situation, the Commander's objectives, and any constraints on the use of resources (e.g. Rules of Engagement, EMCON plan, etc). Thus the basic decision making process always operates within an externally defined frame and is not capable of instituting a transfer to some alternate frame.

Situation Evaluation

The input to the Situation Evaluation element from the environment is taken to consist of I noisy signals, x_i where $i \in I$. The application of a sampling plan on x_i provides the input in the form of a sequence of noisy measures, $x_i(k)$, which are initially taken to correspond to I orthogonal discrete-time dynamic processes; the individual measures, $x_i(k)$, are vectors of dimensionality which taken their values from the finite alphabet X .

The Situation Evaluation element (figure 1) consists of two components, the first of which is concerned with identifying the population from which $x_i(k)$ is drawn. The probability that $x_i(k)$ is a member of the n th population is taken to be:

$$p(f_n(x)|x_i(k)) = \frac{p(f_n(x))p(x_i(k)|f_n(x))}{p(x_i(k))} \quad (1)$$

where $p(f_n(x))$ is the prior probability that an observation, x , will be drawn from the n th population, $p(x_i(k)|f_n(x))$ is the likelihood of a member of the n th population giving rise to the observation $x_i(k)$, and $p(x_i(k))$ is the probability that $x_i(k)$ has its assigned value in X . This representation relies on initial knowledge for the determination of the prior probabilities and the likelihoods; the manner in which these two terms are handled impacts on the adaption of the decision making process.

The values assigned to the set of prior probabilities reflect the expected state of the environment. It is assumed that in any given period a large proportion of these values approximate to zero; in defining a frame these values are set to zero, restricting the estimation process to a subset of the populations. The remaining priors are assigned initial values $p(f_n(x_0))$. In the case of sequential measures on dynamic processes the prior probabilities may be represented as $p(f_n(x_i(k)))$, where

$$p(f_n(x_i(0))) = p(f_n(x_0)) \quad (2)$$

The evolution of the prior probabilities with the successive measures on i may have different forms:

- i) constant prior probabilities

$$p(f_n(x_i(k))) = p(f_n(x_0)) \quad (3)$$

- ii) recursive prior probabilities

$$p(f_n(x_i(k))) = p(f_n(x_i(k-1))|x_i(k-1)) \quad (4)$$

- iii) adaptive prior probabilities which are based on the preceding posterior probabilities adjusted in the light of systematic discrepancies identified by the innovation processes of the individual population models [9].

The first two alternatives have the advantage of simplicity, but the disadvantage that any initial misestimation of the priors is carried through successive estimates.

associated with each population model. It is assumed that the p.d.f. of each model is defined by the

predicted measure, $\hat{x}_{n1}(k)$, and variance, $V_{n1}(k)$, obtained from the population model, where they are derived from the previous measurement history.

The population models, $f(x)$, provide the means of generating state estimates, $\hat{x}_{n1}(k)$, which take the form of vectors of dimensionality s , where

$$s \leq r \quad (5)$$

The individual population model contains a state model, which provides the means of generating the state prediction, $\hat{x}_{n1}(k)$, on the basis of the previous best state estimate, $\hat{x}_{n1}(k-1)$; in addition it contains a measurement model which is used in the derivation of $\hat{x}_{n1}(k)$ from $\hat{x}_{n1}(k)$ and $x_i(k)$, and the generation of $\hat{x}_{n1}(k)$ from $\hat{x}_{n1}(k)$. The variance, $V_{n1}(k)$ is derived from the plant noise of the state model and the measurement noise of the measurement model.

The final task of the Situation Evaluation element is that of electing the hypotheses (populations and state estimates) which should be discarded. The objectives here are those of controlling the total load on the decision maker, and of retaining sufficient alternatives to permit errors to be identified and corrected. It is assumed that a larger number of hypotheses will be retained within Situation Evaluation than will be output to Response Allocation, and that the latter set is included within the former.

The selection process is based on the assignment of costs L_{n1} (where $n, 1 \leq n \leq N$) to the populations, the definition of the costs forming part of the definition of the frame. The selection process differs from that normally used in signal estimation in that hypotheses are tested for rejection against criteria based on the total number of hypotheses extant. The minimum cost for any hypothesis is determined for each input signal, and hypotheses are ordered with respect to costs relative to these minimum costs; hypotheses are then discarded in order of decreasing relative cost until the required numbers of surviving hypotheses are attained.

Response Allocation

The Response Allocation element (figure 2) contains a range of control models, $h_i(x)$, capable of generating control vectors, $y_i(k)$, on the basis of state estimates, $\hat{x}_{n1}(k)$; additionally it contains the means of selecting the control model to be applied to a particular state estimate. The selection of the

control models is governed by an externally defined goal state, Z_G , the current state estimate, Z_C , and externally defined constraints on the use of the control models. In distinction from the model of Boettcher and Lewis [8], the aggregation of the individual state estimates, \hat{x}_{n1} , to form Z_0 is included in the selection process.

The form of selection process employed depends on the nature of the decision maker's environment and on his modes of interaction with it. It is assumed that the state evolves stochastically at a tempo which is not subject to the decision maker's control; additionally it is assumed that it is not generally possible to determine the complete chains of operations which transform Z_0 to Z_G , and that the decision maker is provided with a set of independent resources capable of acting on his environment.

These general assumptions are extended by the further, alternative, sets of assumptions:

- i) that in the timescale of interest the individual \hat{x}_{n1} are independent, and that each control model, $h_i(z)$, corresponds to a mode of operation of a single resource on a particular population or group of populations;
- ii) that in the timescale of interest the \hat{x}_{n1} are best considered in terms of Z , and that each control model, $h_i(z)$, corresponds to the operation of a number of resources.

These sets of assumptions may be related to the different levels of abstraction in ABSTRIPS [7], [10].

In the selection process the control models are represented as stochastic operators which transform state z_{n1} (the state estimate arising from the operation of $f(x)$ on x_i) into state z_{p1} with probability $p(z_{p1} | h_i(z_{n1}))$; feedback from the control models will be used to update this probability during the implementation of the control. The combination of this probability with $p(z_{n1})$, where

$$p(z_{n1}) = p(f_n(x) | x_i) \quad (6)$$

provides the means of evaluating the probabilities of achieving intermediate states, $p(Z)$, through the use of a particular range of processes. The Z may then be evaluated to determine their distances from Z_G , $S(Z_p)$, to provide pruning criteria [7].




Fig.2 Response Allocation

In practise limitations have to be set on the breadth and depth of the search. It has already been mentioned that the definition of the frame includes the provision of constraints on the selection of individual control models; these appear as weights on the selection, $W(z, p(z))$. Alternative approaches using these weights are goal-directed game tree generation using a simplified state and operator representation [10], the identification and use of appropriate pre-planned strategies [11], or the use of Expert System Techniques. The alternative selected will depend on the nature of the decision maker's environment and task.

The outcome of this selection process is the provision of subsidiary goal states to the individual control models. A standard control theory approach is assumed for the control models and the generation of the control vectors, y_i [3]; each control model is based on a single physical process model.

3. THE ADAPTIVE DECISION MAKER

The model of the basic decision making process shows how the processing load on the decision maker may be controlled through external limitation of the available options and the provision of a single goal state. The extension of the model to accommodate self-adaptive, multi-goal seeking behaviour is achieved through the interaction of a pair of decision making processes (figure 3).

The 'Assessment' process observes the external environment and generates control vectors to modify the state of the environment towards the goal state; in carrying out these activities it is constrained by the frame defined by 'Planning'. The Planning process monitors and directs the operation of Assessment; direction is effected by control of the Assessment frame; a possible additional control mode is that of command over-ride or veto. The relationship between Assessment and Planning differs from that between separate interacting decision makers in that they share a common input from the environment, and that their tempos are coupled.

The operation of Assessment is readily represented in terms of the basic decision making process and will not be considered further. Planning also operates in accordance with the basic goal-seeking model of the

decision making process; its goal is that of selecting the optimum combination of externally defined goal and predefined frame for Assessment in the context of the observations of the environment and of Assessment.

The signal inputs to Planning consist of the measures of the environment and on Situation Evaluation and Response Allocation in Assessment (e.g. $p(f(x)|x, k)$, $Z_0, Z, p(Z), S(Z)$). The control inputs include goal states, frames, and frame parameters; the Planning frame, the related set of goal states and their associated weights are expected to be changed infrequently; the frame parameters are expected to be changed more frequently, corresponding to the downward flow of information on the overall situation [4]. Each control model, $h_i(z)$, corresponds to a signal Assessment frame; hence the control vectors, $y_i(k)$, relate either to the alteration of parameters within a frame, or the definition of a new frame.

4. THE CONTROL PROCESS

The interaction between the decision maker and the command system is represented in terms of the Control Process (figure 4). In this model 'Perception' provides the data collection, fusion and filtering required to generate x ; and 'Execution' represents the operation of decision maker's resources.

There is a considerable body of work on the problem of signal detection and estimation [12], [13]; a factor which is observable in this work is the asymptotic behaviour of $p(f(x)|x, k)$ with sample size, k . This convergence has been addressed by Cramer [14] in terms of a sampling distribution which uses the least squares principle to provide a measure of the deviation between the observed and predicted frequencies on the measurement. In this analysis it is shown that the specified sampling distribution tends towards the Chi-squared distribution as $k \rightarrow \infty$; it is also stated that the sampling distribution approximates to the Chi-squared distribution when the predicted frequencies in each sampling region are > 10 .

The operation of the decision maker within the Control Process may be illustrated by the consideration of the simple optimal control problem of minimising the objective function

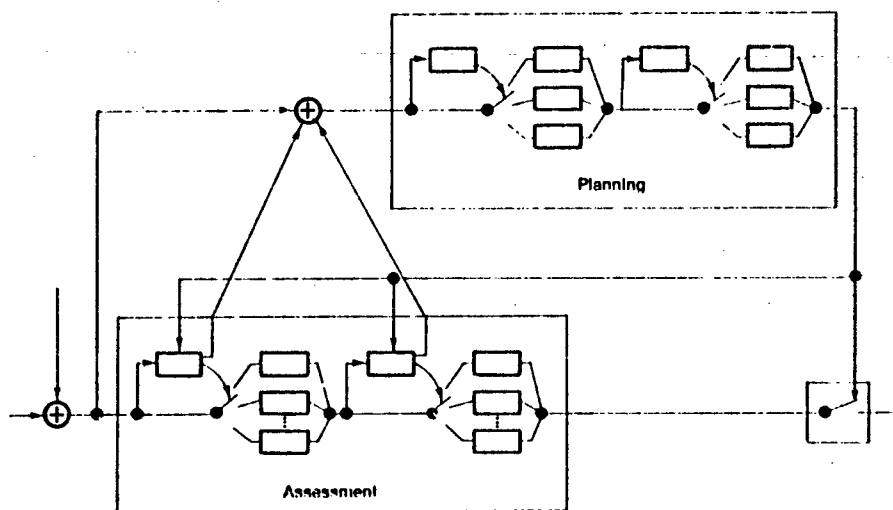


Fig. 3 Adaptive Decision Maker

The operation of Assessment is readily represented in terms of the basic decision making process and will not be considered further. Planning also operates in accordance with the basic goal-seeking model of the decision making process; its goal is that of selecting the optimum combination of externally defined goal and predefined frame for Assessment in the context of the observations of the environment and of Assessment.

The signal inputs to Planning consist of the measures of the environment and on Situation Evaluation and Response Allocation in Assessment (e.g. $p(f(x)|x, (k))$, $Z_0, Z_1, p(Z_1), S(Z_1)$). The control inputs include goal states, frames, and frame parameters; the Planning frame, the related set of goal states and their associated weights are expected to be changed infrequently; the frame parameters are expected to be changed more frequently, corresponding to the downward flow of information on the overall situation [4]. Each control model, $h(z)$, corresponds to a signal Assessment frame; hence the control vectors, $y(k)$, relate either to the alteration of parameters within a frame, or the definition of a new frame.

4. THE CONTROL PROCESS

The interaction between the decision maker and the command system is represented in terms of the Control Process (figure 4). In this model 'Perception' provides the data collection, fusion and filtering required to generate x ; and 'Execution' represents the operation of decision maker's resources.

There is a considerable body of work on the problem of signal detection and estimation [12], [13]; a factor which is observable in this work is the asymptotic behaviour of $p(f(x)|x, (k))$ with sample size, k . This convergence has been addressed by Cramer [14] in terms of a sampling distribution which uses the least squares principle to provide a measure of the deviation between the observed and predicted frequencies on the measurement. In this analysis it is shown that the specified sampling distribution tends towards the Chi-squared distribution as $k \rightarrow \infty$; it is also stated that the sampling distribution approximates to the Chi-squared distribution when the predicted frequencies in each sampling region are > 10 .

The operation of the decision maker within the Control Process may be illustrated by the consideration of the simple optimal control problem of minimising the objective function

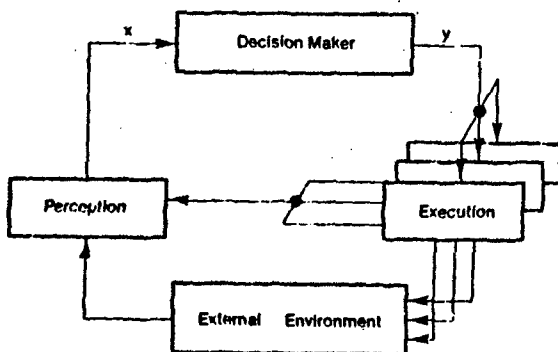


Fig. 4 Control Process

$$\theta(z, y) = \int_0^T [z^2(t) + y^2(t)] dt \quad (7)$$

with alternative process models, one of which accommodates the presence of a hidden variable; using the Euler-Lagrange method, general solutions for (7) are obtained for the alternative process models, these solutions correspond to individual $h(z)$. Figure 5 shows the evolution of the state variable resulting from this control, the lower curve is the optimal control path and the upper curve is the control path resulting from the effect of the hidden variable. The figure also shows three basic periodicities in the operation of the Control Process:

- i) the optimisation period;
- ii) the optimisation cycle corresponding to the selection of $h(z)$ and the updating of parameters in $h(z)$;
- iii) the control cycle corresponding to the tempo of $z_1(k)$ and consequently the tempo of operation of individual $h(z)$.

Application of these periodicities to the adaptive decision maker results in the following set of relationships:

Assessment	Planning
Control Cycle	-
Optimisation Cycle	Control Cycle
Optimisation Period	Optimisation Cycle
-	Optimisation Period

In the simplest case, using Cramer's conclusions for sampling distributions with a single degree of freedom, it appears that optimal performance will be obtained by an adaptive decision maker when the tempo ratio between the Assessment control cycle and the Planning optimisation period is of the order of 1000:1.

This approach may be extended to address hierarchical and lateral interactions within a Command System; the general concepts involved have been discussed in previous papers, [2] and [4].

5. CONCLUSIONS

The application of a control and estimation theoretic approach to the modelling of the decision making process provides a means of identifying different modes of interaction between the decision making process and its environment. It also shows how external knowledge may be used to limit the scope of the decision making process in accordance with some higher level appreciation of the situation; this limits the range of alternatives to be pursued, both in assessing the situation and in selecting appropriate responses, with a consequent reduction of the processing load.

A case which is of particular interest is that of a decision maker capable of self-adaption and of choosing between alternative objectives in response to changes in the environment. A model of an adaptive decision maker is developed from the model of the basic decision making process by the interconnection of two such




Fig. 5

processes; this provides the means of representing complex decision making behaviour in terms of relatively simple elements, and in a manner which reduces processing requirements.

The incorporation of the decision maker within the Control Process provides a basis for analysing the interactions between decision makers within a Command System in terms of tempo of operation and its impact on decision maker performance. This has obvious implications in the area of the development and evaluation of Command System architectures.

6. REFERENCES

- 1) M. Athans, 'the Expert Team of Experts approach to Command and Control Organisations', Proc. 5th MIT/ONR Workshop on C³ Systems, Vol 4, LIDS-R-1159, LIDS, MIT, Cambridge MA, October 1981;
- 2) P.D. Morgan, 'a Total System approach to Command System Design', Proc. 5th MIT/ONR Workshop on C² Systems, LIDS-R-1267, LIDS, MIT, Cambridge MA, December 1982;
- 3) R.F. Kalman, 'On the General Theory of Control Systems', Proc. IFAC Moscow Congress, Vol 1, Butterworth, Inc., Washington D.C. 1960;
- 4) P.D. Morgan, 'Information Management for Command and Control', Proc 5th MIT/ONR Workshop on C³ Systems, LIDS-R-1267, LIDS, Cambridge MA, December 1982;
- 5) J. Klir and M. Valach, 'Cybernetic Modeling', Prague, 1965;
- 6) M.E. Mesarovic, D. Macko and Y. Takahara, 'Theory of Hierarchical, Multilevel Systems', Academic Press, New York, 1970;
- 7) A. Barr and E.A. Feigenbaum, 'The Handbook of Artificial Intelligence', Vol 1, Pitman Books Ltd, London, 1981;
- 8) K.L. Boettcher and A.H. Levis, 'Modeling the Interacting Decision Maker with Bound. Rationality', IEEE Trans. Sys. Man & Cyb., Vol SMC-12, No 3, May/June 1982;
- 9) A.H. Jazwinski, 'Stochastic Processes and Filtering Theory', Academic Press, New York, 1970;
- 10) E.D. Sacerdoti, 'Planning in a hierarchy of abstraction spaces', Artificial Intelligence, 5, 1974;
- 11) D. Michie and I. Bratko, 'Advice Table Representations of Chess End game Knowledge', Proc. AISR/GI Conf. on Artificial Intelligence, 1978;
- 12) S.O. Rice, 'Mathematical Analysis of Random Noise', BSTJ, Vol 24, 1945;
- 13) W.W. Peterson et al, 'The Theory of Signal Detectability', IRE Trans. Inform. Theory, PGIT-4, Sept 1954;
- 14) H. Cramer, 'Mathematical Methods of Statistics', Part 3, Princeton University Press, 1946.

INFORMATION THEORETIC MODELS OF MEMORY IN HUMAN DECISIONMAKING MODELS

Susan A. Hall

RCA Astroelectronics
Princeton, New Jersey

Alexander H. Levis

Massachusetts Institute of Technology
Cambridge, Massachusetts

ABSTRACT

Models of memory and information storage useful in the modeling and analysis of decisionmaking with bounded rationality are discussed. An information theoretic model of permanent memory is presented for describing the accessing of stored information by the algorithms within the human decisionmaker model. It is then applied to the study of the performance - workload characteristic of a decisionmaker performing a dual task.

1. INTRODUCTION

Information theory was first developed as an application in communication theory [1]. But, as Khinchin [2] showed, it is also a valid mathematical theory in its own right, and it is useful for applications in many disciplines, including the modeling of simple human decisionmaking processes [3] and the analysis of information-processing systems. Laming [4] observed, however, that the human decisionmaker does not act like a memoryless communications channel, and, in fact, the purpose of most decisionmaking systems is quite other than to reproduce faithfully at the output what was given to the system as input. In accordance with this observation, a two-stage information theoretic model of the decisionmaking process has been developed [5], [6], [7] which includes internal variables and algorithms between the input and the output. However, the model is memoryless; that is, it is unable to recognize any statistical dependence that might exist in the input or access internal or external data bases. This is a simplifying but very limiting assumption: certainly many organizations receive a variety of inputs related to the same situation, and many of these are statistically dependent on one another. Sen and Drenick [8] recognized the need for adding memory to models of decisionmaking systems. They modeled the human decisionmaker as an adaptive channel, i.e., a channel whose input may depend on present and past inputs. With this addition of memory, they achieved results which, in some experimental situations, reflect observed behavior. However, they have made no attempt to model explicitly the various types of memory that may be found in a decisionmaking system.

Several models of memory have been developed [9]. Buffer storage allows the decisionmaking system to process sequential statistically dependent inputs simultaneously. Permanent memory provides decisionmaking systems with information which is not updated as a result of internal processing, while temporary memory allows for the updating of the stored information. All three have been analyzed [9]; however, emphasis is placed in this paper on a model of permanent memory and its use in the analysis of a model of a human decisionmaker faced with a dual task.

In complex situations when a limited amount of time is available for the decisionmaking process, the decisionmaker may be better modeled as being boundedly rational, i.e., constrained in his abilities to formulate actions and foresee consequences. Rather than always being able to make the optimal decision, a decisionmaker with bounded rationality may satisfy, that is, may seek to satisfy some set of minimal criteria in making a decision [10].

The model of a decisionmaker with bounded rationality [5], [6], [7], shown in Figure 1, consists of two stages: the situation assessment (SA) and the response selection (RS) ones.




Figure 1. Model of decisionmaking process with performance evaluation mechanism

In the SA stage, one of U algorithms is selected via the variable u to evaluate the input and "hypothesize about its origin." The output of the SA stage, z , could be an estimate of the actual signal given the observed input, or some other statistic of the input, or even the entire input itself. The variable z is then given to the response selection stage (RS), and one of V algorithms is chosen via the variable v , to process the evaluated input into an appropriate response. Both sets of algorithms are assumed to be deterministic, so that, given an input x , and the values of u and v , the output y may be exactly determined. Bounded rationality is modeled by requiring that the total rate of activity of the system, where total rate of activity is a well-defined information theoretic quantity, be less than some maximum value, which is specific to a given decisionmaking system.

The performance of the decisionmaker is evaluated as shown in Figure 1. The actual input x' is corrupted by noise, n , so that the system receives $x = x' + n$, a noisy version of the input. This noise could range from representing actual interference with a message sent to the decisionmaker along standard communications channels, to representing the decisionmaker's inability to observe perfectly, or

obtain perfect information pertaining to his environment. The mapping $L(x')$ yields y' , which is defined as the ideal response to the actual input x' ; then y' is compared to the output of the system, y . The performance measure of the system is J , the expectation of $d(y, y')$, where the latter is the cost of deciding y when y' is the desired response. In the context of this model, then, a satisficing decisionmaker must choose a decision strategy, i.e., two probability distributions on u and v , that result in $J \leq \bar{J}$, where \bar{J} is the maximum cost that can be tolerated.

The modeling is developed in the analytic context of N -dimensional information theory. There are two quantities of primary interest. The first of these is entropy: given a variable x , which is an element of the alphabet X , and occurs with probability $p(x)$, the entropy of x , $H(x)$, is defined to be

$$H(x) = - \sum_x p(x) \log p(x) \quad (1.1)$$

and is measured in bits when the base of the logarithm is two. Entropy is also known as the average information or uncertainty in x , where information does not refer to the content of the variable x , but rather to the average amount by which knowledge of x reduces the uncertainty about it. The other quantity of interest is average mutual information or transmission: given two variables x and y , elements of the alphabets X and Y , and given $p(x)$, $p(y)$, and $p(x|y)$ (the conditional probability of x , given the value of y), the transmission between x and y , $T(x:y)$ is defined to be

$$T(x:y) = H(x) - H_y(x) \quad (1.2)$$

where

$$H_y(x) = - \sum_y p(y) \sum_x p(x|y) \log p(x|y) \quad (1.3)$$

is the conditional uncertainty in the variable x , given full knowledge of the value of the variable y .

McGill [11] extended this basic two-variable input-output theory to N dimensions by extending Eq. (1.2):

$$T(x_1, x_2, \dots, x_N) = \sum_{i=1}^N H(x_i) - H(x_1, x_2, \dots, x_N) \quad (1.4)$$

For the modeling of memory and of sequential inputs which are dependent on each other, the use of the entropy rate, $\bar{H}(x)$, which describes the average entropy of x per unit time, is appropriate:

$$\bar{H}(x) = \lim_{m \rightarrow \infty} \frac{1}{m} H[x(t), x(t+1), \dots, x(t+m-1)] \quad (1.5)$$

Transmission rates, $\bar{T}(x:y)$, are defined exactly like transmission, but using entropy rates in the definition rather than entropies.

Conant's Partition Law of Information Rates (PLIR) [12] is defined for a system with $N-1$

internal variables, w_1 through w_{N-1} , and an output variable, y , also called w_N . The PLIR states:

$$\sum_{i=1}^N \bar{H}(w_i) = \bar{T}(x:y) + \bar{T}_y(x:w_1, w_2, \dots, w_{N-1}) + \bar{T}(w_1, w_2, \dots, w_{N-1}:y) + \bar{H}_x(w_1, w_2, \dots, w_{N-1}, y) \quad (1.6)$$

and is easily derived using information theoretic identities. The left-hand side of Eq. (1.6) refers to the total rate of activity of the system, also designated \bar{G} . Each of the quantities on the right-hand side has its own interpretation. The first term, $\bar{T}(x:y)$, is called the throughput rate of the system and is designated \bar{G}_t . It measures the amount by which the output of the system is related to the input.

$$\bar{T}_y(x:w_1, w_2, \dots, w_{N-1}) = \bar{T}(x:w_1, w_2, \dots, w_{N-1}:y) - \bar{T}(x:y) \quad (1.7)$$

is called the blockage rate of the system and designated \bar{G}_b . Blockage may be thought of as the amount of information in the input to the system that is not included in the output. The third term, $\bar{T}(w_1, w_2, \dots, w_{N-1}:y)$, is called the coordination rate of the system and designated \bar{G}_c . It is the N -dimensional transmission of the system; i.e., the amount by which all of the internal variables in the system constrain each other. The last term, $\bar{H}_x(w_1, w_2, \dots, w_{N-1}, y)$ designated \bar{G}_n represents the uncertainty that remains in the system variables when the input is completely known. This noise should not be construed to be necessarily undesirable as it is in communications theory: it may also be thought of as internally-generated information, information supplied by the system to supplement the input and facilitate the decisionmaking process. The PLIR may be abbreviated:

$$\bar{G} = \bar{G}_t + \bar{G}_b + \bar{G}_c + \bar{G}_n \quad (1.8)$$

The bounded rationality constraint is expressed by postulating the existence of a maximum rate of information-processing, or a maximum rate of total activity, \bar{G}_{max} , at which a given decisionmaking system can operate without overload. Note that the addition of memory to the decisionmaking model increases the total number of variables in the system and may, therefore, restrict the strategies that may be used to those with lower activity or workload. However, executing a task with memory may result in a better performance than that achievable in a system without memory.

In the next section, a model of memory is presented. In the third section, the model is used to study the performance - workload characteristics of a decisionmaker assigned with the execution of two concurrent tasks — the dual task problem.

2.0 PERMANENT AND TEMPORARY MEMORY

2.1 Introduction

Memory is assumed to consist of both permanent and temporary stores of information which may be drawn upon by the algorithms in the situation assessment and the response selection stages during the decisionmaking process. Permanent memory is defined as to contain values which are constant; that is, they may not be revised or appended by the algorithms but access them. Temporary storage contains values which may be revised by the algorithms; for example, a discrete Kalman filter algorithm would include temporary storage of the best estimate of the present state of the process, to be used in the next iteration of the algorithm. Temporary memory has the effect of adding memory to the algorithms themselves; with temporary memory available, the algorithms can remember values from one iteration to the next. The division of memory into permanent and temporary bears a strong resemblance to the division of memory that is made in the cognitive sciences, into long-term and short-term memory [13].

A third type of memory, called sensory memory, is also hypothesized by psychologists. Information from the environment is stored in sensory memory before it undergoes any processing; sensory memory might therefore be compared to a buffer storage model. The latter allows the simultaneous processing of sequential statistically dependent inputs. Several different models have been developed [9] that depend on the class of inputs that the system receives. Shift register buffers provide the storage rule necessary to process input from a general Markov source. Fixed-length string buffers are a suitable model for the type of storage found in machines. Variable-length string buffers are appropriate models for some types of human sensory memory. Shift register buffers are simple, but add a great deal of activity to the system and result in redundant processing. Fixed-length buffers do not suffer from these deficiencies, but introduce a substantial delay which is proportional to the length of the string. Variable-length string buffers have smaller average delay than fixed-length ones, but increase the overall activity because of their relative complexity.

The model of permanent memory presented in this paper is similar to long-term memory, in that information is stored indefinitely and is accessible by information processing mechanisms. It is different in that new information is being added continuously to long-term memory; the permanent memory model in this paper provides no mechanism for this addition. Second, information may be lost from long-term memory; this permanent memory model does not have a forgetting mechanism. These differences are noted to indicate that, although similarities exist between the model of memory presented here and that found in the cognitive sciences, permanent memory is not intended to be a model of long-term memory per se.

Permanent memory may be accessed by both the situation assessment and the response selection stages. However, in this paper, consideration will be limited to the situation assessment stage. The relationships derived are the same as they would be for the response selection stage, since the two halves of the decisionmaking process are structurally identical.

It is quite possible that a decisionmaking system may contain both buffer storage and permanent and temporary units. However, in order to simplify the

presentation, the assumption is made that the decisionmaking system contains no buffer storage. This situation is relaxed easily [9].

The general memory unit applicable to both permanent and temporary memory, is shown in Figure 2. It consists of M variables, d_1 through d_M , as well as an input M -vector, D_I , and an output M -vector, D_O . Note that because permanent memory may not be revised, its model will not contain the input vector D_I .

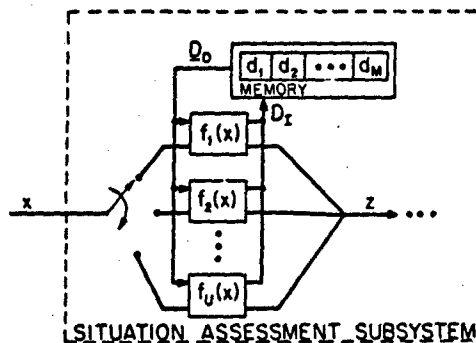


Figure 2. Model of SA subsystem with memory

2.2 Permanent Memory

It might seem at first that the addition of permanent memory to a decisionmaking system might have no effect at all on the total information theoretic rate of activity of the system; if the values of d_k for $k=1,2,\dots,M$ do not change over time, then

$$\bar{H}(d_k) = 0 \quad k = 1,2,\dots,M \quad (2.1)$$

Since total activity is just the sum of the entropies of the individual variables in the system, it appears that the addition of M deterministic variables to a system should have no effect on its total activity. However, the problem is actually more complex. In order to demonstrate the types of changes that occur when permanent memory is added to the model, a particularly simple example will be analyzed.

Let the permanent memory unit consist of one variable, d_1 , which may be accessed by one SA algorithm, f_1 , as shown in Figure 3.

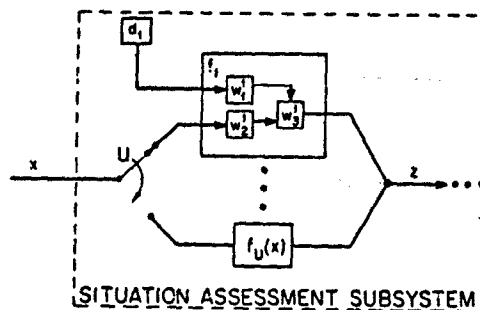


Figure 3. Example of SA subsystem with permanent memory

Algorithm f_1 provides the average value of the two components of a vector input. Whenever a specific algorithm is accessed by the decisionmaking system, the variables of that algorithm are defined to be active, and take values according to some probability distribution which is a function of the input. When

that algorithm is not accessed, its variables are defined to be inactive; i.e., they assume some fixed value, say ϕ , which they may not assume when they are active.

Now consider algorithm f_1 , which provides the mean value of the two components of a vector input. There are two similar ways of implementing this algorithm. The first does not access permanent memory, although there is some implicit memory in the algorithm itself:

$$\begin{aligned} w_1 &= 2; & w_2 &= x_1 + x_2 \\ w_3 &= w_1/w_2; & z &= w_3 \end{aligned} \quad (2.2)$$

The second does access permanent memory:

$$d_1 = 2;$$

defined outside the algorithm f_1

$$\begin{aligned} w_1 &= d_1; & w_2 &= x_1 + x_2 \\ w_3 &= w_2/w_1; & z &= w_3 \end{aligned} \quad (2.3)$$

In the second example, the variables w_1, w_2, w_3 are inactive when the algorithm is not accessed. In the first example, however, w_1 must retain the value of 2 throughout, since no means have been provided to reinitialize its value each time the algorithm is accessed. It is now possible to compare the levels of activity of the system with the permanent memory unit and that without. First consider realization A. The throughput, blockage, and noise rates of the SA subsystem are given by [9]:

$$\begin{aligned} \bar{G}_t &= \bar{H}(z) - \bar{H}_z(z) \\ \bar{G}_b &= \bar{I}_z(z; w_1, \dots, w_{aU}) \\ \bar{G}_n &= \bar{H}(u) \end{aligned} \quad (2.4)$$

With inputs arriving once every second, the coordination rate is found as follows:

$$\bar{G}_c = \sum_{i=1}^U \sum_{j=1}^{a_i} \bar{H}(w_j^i) + \bar{H}(u) + \bar{H}(z) - \bar{H}(u, w, z) \quad (2.5)$$

Here, w_j^i represents the j -th variable of algorithm i ; and W represents the entire set of w_j^i in the SA subsystem. Finally, a_i is the number of internal variables of algorithm i which are active or inactive according to the value of u . Equation (2.5) reduces to:

$$\begin{aligned} \bar{G}_c &= \sum_{i=1}^U a_i \bar{H}[p(u)=i] \\ &+ \sum_{i=1}^U \sum_{j=1}^{a_i} \bar{H}_u(w_j^i) - \bar{H}_u(W) + \bar{H}(z) \end{aligned} \quad (2.6)$$

The symbol \bar{H} denotes the binary entropy of its argument, given by:

$$\bar{H}(p) = -p \log_2 p - (1-p) \log_2 (1-p), \quad 0 \leq p \leq 1 \quad (2.7)$$

The same quantities may be calculated for realization B. The rates of throughput, blockage and noise are not effected by the small structural difference in algorithm f_{1B} ; the rate of coordination does change. Consider Eq. (2.6): a_{1B} is now equal to 3, because w_1 is now active when $u=1$ and inactive otherwise. Therefore, the first term of Eq. (2.6) is increased by the amount $\bar{H}[p(u=1)]$. The second and third terms remain the same, even though there is now some uncertainty associated with the value of w_1 . Knowledge of the value of u resolves that uncertainty, so that

$$\bar{H}_u(w_1) = 0 \quad (2.8)$$

Similarly, $\bar{H}_u(W)$ is unchanged. Only the structure of the algorithm has been changed, so the output remains the same, and the last term of Eq. (2.6) is unchanged. Therefore, the addition of one unit of permanent memory to the SA subsystem provides a total increase in activity of $\bar{H}[p(u=1)]$. In general, if algorithm i directly accesses β_i values from permanent memory, and no other changes are made in the algorithms, then the incremental activity of the system, $\Delta \bar{G}$, is given by

$$\Delta \bar{G} = \sum_{i=1}^U \beta_i \bar{H}[p(u=i)] \quad (2.9)$$

3.0 THE DUAL-TASK PROBLEM

3.1 Introduction

It has been observed that if a person must execute two tasks by switching between them, his level of performance may be different than when he is allowed to confine himself to one task [13], [14], even if the arrival rate for individual tasks is the same for both cases. If there is some synergy between the two tasks—that is, if the two tasks are related and executing one actually helps the execution of the other—then performance may improve. If, on the other hand, the two tasks are dissimilar or simply do not reinforce each other, performance may decline from what it was in the single-task case. It is this latter phenomenon that will be explored in this section.

There are numerous possible ways in which the dual-task problem might be modeled. For example, if the two tasks to be performed are assumed to be so different from each other that they demand different sets of algorithms, then a pre-processor may be required for the system. The pre-processor determines which type of task each input represents and then allows access to a set of decisionmaking algorithms appropriate to that task. Of course, the activity of the pre-processor increases overall system activity and may, therefore, lower performance. On the other hand, if the two tasks to be performed are assumed to be similar but non-synergistic, they may be able to use the same basic sets of algorithms, as long as these algorithms are adaptable to each task through two different sets of parameter values stored in permanent memory. Notice that there is an implicit need for a pre-processor in this problem, since the algorithms must have some way of knowing which type of

It has been received in order to determine which of values stored in memory to access. Overall activity is increased in this formulation as well by necessity of switching between sets of information. An example of this second problem is a switchboard operator who has to process both incoming and outgoing calls; although the tasks are the same basic action, they differ with respect to the information required to execute the task. The second problem is addressed in this paper; first one will be presented at a later time.

In order to simplify this problem, several assumptions will be made. First, to circumvent the need for a pre-processor, it is assumed that there are separate inputs to the system, x_A and x_B , which members of disjoint alphabets, X_A and X_B . Only one of these inputs is active at any given time: if x_A is active, task A must be performed, and if x_B is active, task B must be performed. Inputs arrive at the system once every second, and there is a known probability TD (representing the task division) that x_A will be active at any given time.

If inputs are not synergistic, then they are assumed to be statistically independent. If x is generated independently every τ seconds, then

$$\bar{H}(x) = \frac{1}{\tau} H(x) \quad (3.1)$$

Therefore, in the results which follow, entropies per second rather than entropy rates will be used. Activities are denoted by G in place of \bar{G} for activity rates. The units for G are bits per symbol (as opposed to bits per second for \bar{G}). Note that for the problem with synergy between tasks, the assumption of dependence between sequential inputs would be appropriate; a buffer storage model would be added and activity rates would be used in the analysis.

The basic model for the problem is shown in Figure 4. The variable u , which acts independently of the input x , controls which of two situation assessment (SA) algorithms, f_1 and f_2 , will be used. The decision strategy for a system such as this may then be defined by the probability δ that u is equal to 1. For simplicity, it is assumed that the purpose of both tasks is merely to assess the situation, so that z is the output (no RS stage). The variables s_1 and s_2 are represented as switches internal to the algorithms only so that their function may be highlighted. Figure 4 does not explicitly depict the mechanism by which s_1 and s_2 take their values, but only that they are dependent on the value of u and on the values of x_A and x_B .

Specifically, they take values as follows:

$$\begin{aligned} s_1 &= A \quad \text{if } u = 1, \quad x_A \neq \phi, \quad x_B = \phi \\ s_1 &= B \quad \text{if } u = 1, \quad x_A = \phi, \quad x_B \neq \phi \quad i = 1, 2 \\ s_i &= \phi \quad \text{if } u \neq 1 \end{aligned} \quad (3.2)$$

In addition to s_1 and s_2 , algorithms f_1 and f_2 contain a_1 and a_2 internal variables. Finally, D_A and D_B are the two sets of information or data needed by the algorithms to process input from X_A and X_B , respectively. It is assumed that both algorithms use all of the information in D_A when performing task A and all of D_B for task B.

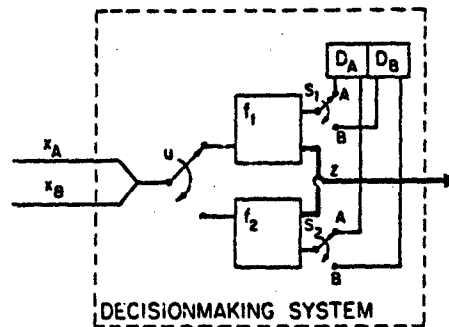


Figure 4. Model of decisionmaker with permanent memory for Dual-Task problem

3.2 Information Theoretic Analysis

In order to measure if a change in performance level occurs between the single-task and dual-task situations, a performance index is required. The index J that is used is the probability of error. In terms of the quantities defined in section 1 and depicted in Figure 1 (with the output of the system now equal to z),

$$d(z, z') = \begin{cases} 1 & \text{if } z \neq z' \\ 0 & \text{if } z = z' \end{cases} \quad (3.3)$$

and therefore J , the expectation of $d(z, z')$ is

$$J = \sum_z p(z) d(z, z') = \text{prob}(z \neq z') \quad (3.4)$$

Because two distinct tasks are being performed, J_A is defined as the probability of error in executing a type A task; and J_B as the probability of error in a type B task. More precisely,

$$J_i = \text{prob}(z \neq z' | x \in X_i) \quad i = A, B \quad (3.5)$$

Note that these quantities are independent of the task division TD, the probability that $x \in X_A$, but will be dependent in general on the decision strategy δ . In fact, if it is known how the system performs when pure strategies are employed (either u is 1 with probability 1, and algorithm f_1 is always used, or u is 2 and algorithm f_2 is always used), the performance of the system under the mixed strategy δ (algorithm f_1 is used with probability δ) is simply a convex combination of the performances using pure strategies [5], i.e.,

$$\begin{aligned} J_i(\delta) &= \delta(J_i | u=1) + (1-\delta)(J_i | u=2) \\ i &= A, B; \quad 0 \leq \delta \leq 1 \end{aligned} \quad (3.6)$$

With this definition of task performance, it is also possible to define an overall performance index for the system:

$$J(\delta) = (TD)J_A(\delta) + (1-TD)J_B(\delta) \quad (3.7)$$

If errors on one task are more detrimental than errors on the other, then weighting coefficients may be introduced on the right-hand side of Eq. (3.7).

The activity of the system will change both as a function of the decision strategy δ , and as a function of the task division TD. In fact, G , the total activity of the system, is convex both in δ (with a fixed TD) and in TD (with δ fixed). The convexity of G in δ has already been shown [5]; the convexity of G in TD will be demonstrated.

Assume that only task A is being performed. Note that under this assumption, the need for variables s_1 and s_2 disappears; the algorithms may be directly connected to data base D_A . In this case, the levels of activity for a decisionmaker with two SA algorithms f_1 and f_2 , containing a_1 and a_2 internal variables, respectively, and a decision strategy δ , are given by:




Figure 5. Representation of G vs. J for binary variation of pure strategies

$$\begin{aligned} G_t + G_b &= H(x) \\ G_n + H(u) &= H(\delta) \\ G_c &= (a_1 + a_2)H(\delta) + H(x) + \sum_{i=1}^2 p(u=i)g_c^i \end{aligned} \quad (3.8)$$

The quantity H is the entropy of a binary variable; g_c^1 and g_c^2 are defined to be the internal coordinations of algorithms f_1 and f_2 , respectively, where internal coordination is defined as

$$g_c^i = \frac{1}{\text{prob}(u=i)} \left[\sum_{j=1}^{a_i} H(w_j^i | u=i) - H(W^i | u=i) \right] \quad (3.9)$$

and W^i represents the set of all of the variables of algorithm i . The total activity of the system is then the sum of the quantities given in Eq. (3.8):

$$\begin{aligned} G &= H(x) + H(z) + (a_1 + a_2 + 1)H(\delta) \\ &+ \sum_{i=1}^2 p(u=i)g_c^i \end{aligned} \quad (3.10)$$

Note that all of the above quantities are conditional on task A being performed: e.g., g_c^i could be written $(g_c^i | x \in X_A)$. It has been shown [5] that G is convex in the decision strategy, i.e.:

$$G(\delta) \geq (\delta)(G|u=1) + (1-\delta)(G|u=2) \quad (3.11)$$

Therefore, using Eqs. (3.6) and (3.11), G may be found parametrically as a function of J for the single-task problem, as shown in Figure 5.

The dual-task problem requires the variables s_1 and s_2 to be included in the model. It is still the case that

$$\begin{aligned} G_t + G_b &= H(x) \\ G_n &= H(u) = H(\delta) \end{aligned} \quad (3.12)$$

Let d_k represent a single variable of the permanent memory unit, let D represent all the permanent memory (both D_A and D_B), and let W represent all of the internal variables of both algorithms f_1 and f_2 ; then the coordination of the system performing two tasks is given by:

$$\begin{aligned} G_c &= \sum_{i=1}^2 \sum_{j=1}^{a_i} H(w_j^i) + H(s_1) + H(s_2) + H(u) + H(z) \\ &+ \sum_k H(d_k) - H(W, s_1, s_2, u, z, D) \end{aligned} \quad (3.13)$$

After much manipulation, Eq. (3.13) may be reduced to

$$\begin{aligned} G_c &= (a_1 + a_2 + 2)H(\delta) \\ &+ \sum_{i=1}^2 p(u=i) [(TD)(g_c^i | x \in X_A) + (1-TD)(g_c^i | x \in X_B)] \\ &+ \sum_{i=1}^2 \sum_{j=1}^{a_i} T(w_j^i : s_i | u=i) + H(z) \end{aligned} \quad (3.14)$$

and the total activity for the system performing two tasks is given by:

$$\begin{aligned} G &= H(x) + H(z) + (a_1 + a_2 + 3)H(\delta) \\ &+ \sum_{i=1}^2 p(u=i) [(TD)(g_c^i | x \in X_A) + (1-TD)(g_c^i | x \in X_B)] \\ &+ \sum_{i=1}^2 \sum_{j=1}^{a_i} T(w_j^i : s_i | u=i) \end{aligned} \quad (3.15)$$

There are now two additional $H(\delta)$ terms; these are due to the presence of the two additional system variables, s_1 and s_2 . The internal coordination term is now a convex combination of the internal coordinations found when only task A or B is performed. Finally, the last term of Eq. (3.15), which does not even appear in Eq. (3.10):

$$T(s_j^i: s_i | u=i) = H(s_i | u=i) - H_j(s_i | u=i) \quad (3.16)$$

This may be interpreted as the amount of information transmitted between s_i and w_j^i , given that algorithm i is being used; i.e., it is the extent to which variable w_j^i reflects which task is being performed. Since

$$H(s_i | u=i) = p(u=i)H(TD) \quad (3.17)$$

then

$$0 \leq T(w_j^i: s_i | u=i) \leq p(u=i)H(TD) \quad (3.18)$$

It will now be shown that for a fixed value of δ , $0 \leq \delta \leq 1$, G is convex in the task division, i.e.,

$$G(TD) \geq (TD)(G|_{x \in X_A}) + (1-TD)(G|_{x \in X_B}) \quad 0 \leq TD \leq 1 \quad (3.19)$$

The right-hand side (RHS) of (3.19) may be found using (3.10):

$$\begin{aligned} \text{RHS} &= (TD)[H^A(x) + H^A(z) + \alpha_1 + \alpha_2 + 1)H(\delta)] \\ &\quad + \sum_{i=1}^2 p(u=i)(g_c^i |_{x \in X_A}) \\ &\quad + (1-TD)[H^B(x) + H^B(z) + \alpha_1 + \alpha_2 + 1)H(\delta)] \\ &\quad + \sum_{i=1}^2 p(u=i)(g_c^i |_{x \in X_B}) \end{aligned} \quad (3.20)$$

Here, $H^j(x)$ and $H^j(z)$ are the entropies of x and z which occur when only a single task is executed. The probability distributions for x and z in the dual-task case are a convex combination of those for the single-task cases,

$$\begin{aligned} p(x) &= (TD)p(x|_{x \in X_A}) + (1-TD)p(x|_{x \in X_B}) \\ p(z) &= (TD)p(z|_{z \in Z_A}) + (1-TD)p(z|_{z \in Z_B}) \end{aligned} \quad 0 \leq TD \leq 1 \quad (3.21)$$

When a probability distribution is the convex combination of two others, as in Eq. (3.21), then [16]:

$$\begin{aligned} H(x) &\geq (TD)H^A(x) + (1-TD)H^B(x) \\ H(z) &\geq (TD)H^A(z) + (1-TD)H^B(z) \end{aligned} \quad 0 \leq TD \leq 1 \quad (3.22)$$

and it follows that Eq. (3.20) can be written as:

$$\begin{aligned} \text{RHS} &= [(TD)H^A(x) + (1-TD)H^B(x)] \\ &\quad + [(TD)H^A(z) + (1-TD)H^B(z)] \\ &\quad + [(\alpha_1 + \alpha_2 + 1)H(\delta)] \\ &\quad + \sum_{i=1}^2 p(u=i)[TD(g_c^i |_{x \in X_A}) + (1-TD)(g_c^i |_{x \in X_B})] \end{aligned} \quad (3.23)$$

Now compare Eq. (3.23) to Eq. (3.15), using the results of Eq. (3.22), the fact that $H(\delta) \geq 0$, and the fact that transmissions must also be non-negative. It follows that Eq. (3.19) does indeed hold, and H is convex in the task division. In fact, if a mixed strategy δ is being used ($0 < \delta < 1$), or if any of the internal variables of an algorithm in use reflects which task is being performed, i.e.,

$$T(w_j^i: s_i | u=i) > 0 \quad (3.24)$$

then the inequality of Eq. (3.19) will be strict.

3.3 Effect of Task Division on Performance

To see the effects that this result has on performance, consider a particularly simple example. It is assumed that the single-task activity or workload versus performance curves are identical for task A and task B (this implies that J_A and J_B are the same functions of δ ; see Figure 6a). Now consider the evolution of the G versus J_A curve as TD changes from 0 to 1. It is meaningless to define J_A for the single-task case in which task B is always performed ($TD = 0$), but for very small values of TD , J_A is defined as in Eq. (3.5). To find the G versus J_A curve for $TD \approx 0$, consider Eq. (3.15). Since $H(TD) \approx 0$ for $TD \approx 0$, its last term is small (see Eq. (3.18)). The rest of Eq. (3.15) reduces to Eq. (3.25):

$$G(TD \approx 0) \approx (G|_{x \in X_B}) + 2H(\delta) \quad (3.25)$$

In other words, the G versus J_A curve will be the same as either single-task curve, with the quantity $2H(\delta)$ added on due to the presence of variables s_1 and s_2 (see Figure 6b). As TD increases, G will continue to increase up to some point (because of its convexity in TD), dependent on the value of the last term of Eq. (3.15) and the values of $H(x)$ and $H(z)$. For TD equal to 0.5, the G versus J_A curve will have the general shape shown in Figure 6c. Finally, G will decrease until $TD \approx 1$ and G versus J_A is again as shown in Figure 6b. For a fixed value of δ then, say $\delta = 0.2$, the workload versus task division curve will be similar to that shown in Figure 6d. The maximum activity need not occur at $TD = 0.5$.

4.0 CONCLUSION

In order to obtain more realistic models of humans carrying out information processing and decisionmaking tasks, it is necessary that memory, whether internal to the decision process, or external in the form of data bases, be modeled. Three classes of models are described: buffer storage, permanent memory and temporary memory. The modeling of permanent memory has been presented and illustrated through its use to the analysis of the performance-workload characteristic of a human decisionmaker executing a dual task.

In order to test experimentally these predictions the model for the dual-task problem as defined here, several criteria must be met. First, the two tasks must be similar enough that the same set of algorithms may be used for both tasks; however, they should be independent enough so that execution of one task does not aid in the execution of the other. Second, it should be necessary to switch between tasks, i.e., two different tasks may not be performed simultaneously. Third, individual tasks should arrive at the same rate in the dual-task test as in the single-task test. Finally, this rate of presentation should be near to the bounded rationality constraint of the decisionmaker, since it is hypothesized that it is this constraint that leads to performance degradation.

5.0 ACKNOWLEDGMENT

This work was carried out at the MIT Laboratory for Information and Decision Systems with support by the U. S. Air Force Office of Scientific Research under contract No. AFOSR-80-0029.

6.0 REFERENCES

- [1] C. E. Shannon and W. Weaver, The Mathematical Theory of Communication, University of Illinois, Urbana, 1949.
- [2] A. I. Khinchin, Mathematical Foundations of Information Theory, Dover Publications, Inc., New York, 1957.
- [3] T. B. Sheridan and W. R. Ferrell, Man-Machine Systems, MIT Press, Cambridge, MA, 1974.
- [4] D. R. J. Laming, Information Theory of Choice-Reaction Times, Academic Press, London and New York, 1968.
- [5] K. L. Boettcher and A. H. Lewis, "Modeling the Interacting Decisionmaker with Bounded Rationality," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-12, No. 3, May/June 1982.
- [6] A. H. Lewis and K. L. Boettcher, "Decisionmaking Organizations with Acyclical Information Structures," IEEE Transactions on Systems, Man, and Cybernetics, SMC-13, May/June 1983.
- [7] A. H. Lewis and K. L. Boettcher, "On Modeling Teams of Interacting Decisionmakers with Bounded Rationality," Proceedings IFAC/IFIP/IFORS/IEA Conference on Analysis, Design, and Evaluation of Man-Machine Systems, VDI/VDE, Duesseldorf, FRG, September 1982; to appear in Automatica, Vol. 19, No. 6, November 1983.




Figure 6. Performance vs. Workload and Task Division

Now consider what happens to performance if the maximum total workload constraint is given by the value marked G_{max} in Figure 6, i.e., the system is required to perform at an activity level $G \leq G_{max}$. In the two single-task cases, the system is unconstrained and may use any strategy $0 \leq \delta \leq 1$. However, for this example, when both tasks arrive with equal probability ($TD = 0.5$), the set of feasible strategies is greatly reduced, and performance is limited to being very poor. Also, the particular strategy of $\delta = 0.2$ may only be used for task divisions close to 0 or 1 (see Figure 6d).

This simple example illustrates some rather general results. The convexity of G in the task division implies that the rate of activity of the system will be greater in the dual-task case than in at least one of the single tasks. If the workload of the two tasks is very disparate, then the opportunity to switch between a very activity-intensive (high workload) task and a very easy one may actually reduce the workload from what it is in the case that only the difficult task is being performed. When the activity levels for the two single-task cases are comparable, though, as in the preceding example, the workload for the dual-task case is greater than that for either of the single-task cases. This increase in workload arises from three basic sources, which may be seen by an examination of Eq. (3.15). First, in the dual-task case, the variable x and in most cases also the variable z will have a larger uncertainty associated with them because of their larger alphabets. Second, the dual-task problem requires that the system have some means of switching between the sets of data stored in permanent memory: the variables x_1 and x_2 provide the mechanism but also increase the uncertainty of the system. Third, the rest of the internal variables may, because of access to different values stored in memory, take on a wider range of values when both tasks must be performed than when only one is performed. If the system performing either task alone is operating near its maximum allowable rate, then requiring the system to switch between the two tasks has the effect of both eliminating the more active decision strategies from the feasible set, and, in the case that the more active strategies also result in better performance, lowering the performance of the system.

- [8] P. Sen and R. F. Drenick, "Information Processing in Man-Machine Systems," Proc. of the Fourth MIT/ONR Workshop on C Systems, LIDS-R-1159, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, October 1981.
- [9] S. A. Hall, "Information Theoretic Models of Storage and Memory," S. M. Thesis, LIDS-TE-1232, Laboratory for Information and Decision Systems, MIT, Cambridge, MA, August 1982.
- [10] J. G. March and E. A. Simon, Organizations, John Wiley and Sons, New York, 1958.
- [11] W. J. McGill, "Multivariate Information Transmission," Psychometrika, Vol. 19, No. 2, June 1954.
- [12] R. C. Conant, "Laws of Information Which Govern System," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-6, No. 4, April 1976.
- [13] R. L. Solso, Cognitive Psychology, Harcourt, Brace, and Jovanovich, Inc., New York, 1979.
- [14] R. T. Kelly, F. L. Greitzer, and R. L. Hershman, "Air Defense: A Computer Game for Research in Human Performance," NPRDC TR 81-15, Navy Personnel Research and Development Center, San Diego, CA, July 1981.
- [15] R. T. Kelly, and F. L. Greitzer, "Effects of Track Load on Decision Performance in Simulated Command and Control Operations," NPRDC TR 82-21, Navy Personnel Research and Development Center, San Diego, CA, January 1982.
- [16] R. G. Gallager, Information Theory and Reliable Communication, John Wiley and Sons, New York, 1968.

AD P002884

HUMAN DECISIONMAKING IN DYNAMIC ENVIRONMENTS WITH INCREASING INFORMATION PROCESSING DEMANDS

Daniel Serfaty, Eric P. Soulsby, and David L. Kleinman

CYBERLAB, Dept. of Electrical Engineering and Computer Science
University of Connecticut
Storrs, CT 06268

ABSTRACT

In order to describe the overall behavior of a man-machine system it has long been recognized that there exists a need to determine the level of "workload" imposed on the human operator in achieving design objectives. One of the objectives of workload related research is to determine areas of high workload, a subset of which entails the domain of information overload. Several mechanisms of adjustment to increasing information processing demands are available to the decision maker; e.g. queueing, filtering, omission, and employing multiple channels. These mechanisms of adjustment have an effect on the human's operative behavior and subsequent strategy. When presented with increasing information processing demands near or in excess of his capacity the decisionmaker exhibits changes in operating methods to avoid crossing an overload threshold. Hence the operative methods of the decisionmaker are efficient from the point of view of performance, while being economical from the point of view of workload.

The Dynamic Decision Model (DDM)_A is a normative-descriptive model that has shown to provide an excellent representation of human information processing and decisionmaking in a dynamic multi-task environment. In the present effort, a sensitivity study was performed on the DDM in an attempt to explore the nature by which various dynamic task attributes affect human performance. In particular, changes in parameters such as the number of concurrent tasks, task velocity, task processing time and task value, were investigated in conjunction with existing notions concerning human operator workload. This global approach not only emphasizes the performance level the human operator may attain but also the tactics and strategies he uses to achieve it.

Results indicate a general agreement with existing workload/performance theories and some inherent human information processing limits are identified. Extensive experimental studies were performed along these lines by using an experimental paradigm which abstracts some feature of a C³ decisionmaking situation. They quantitatively confirm, to a large extent, the DDM analytical predictions of human performance sensitivities. In particular, the data show an interesting feature of human behavior: the ability to adapt to high workload situations by discriminating some alternatives and maintaining performance through some "satisfaction" criterion. A subjective workload rating technique (SWAT) was used to confirm the perceived increase in information processing demands.

I INTRODUCTION

I.1 Background and Motivation of the research

Previous research at the University of Connecticut CYBERLAB was aimed at modeling single human decision-making processes in multiple task environments. The experimental part of this effort focused on developing a canonical decision paradigm through which the deci-

sionmaker's performance and behavior were studied as task parameters were changed [1-3]. This paradigm was motivated, in large part, by the C³ problem of target selection. In this situation, targets of various type move across the display scopes of the human operator, vying for his attention. Each target has a different threat value and processing time requirement. The human, therefore, is faced with the problem of sequencing tasks dynamically so as to maximize the performance of the system.

Fig. 1 shows the fundamental decision loop that is considered in our approach. The human information and decision process involves 1) whether to process a task or gather more information (i.e. monitor); and 2) which one of N tasks to act upon (N is time varying), in order to maximize the system performance. The decision loop is dynamic in nature. As time evolves, tasks of

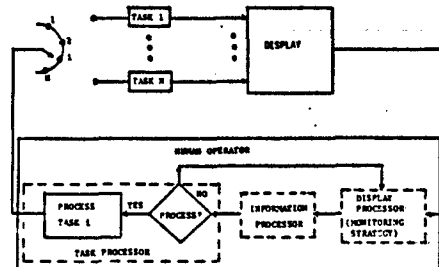


Fig. 1. MULTI-TASK DYNAMIC MONITORING/DECISION LOOP

different value, duration (processing time) and opportunity window demand the operator's attention, while others depart. The opportunity windows shrink with time as the tasks approach their deadlines.

In a supervisory, decisionmaking situation the human operator must process information presented en route to choosing an appropriate course of action. With regard to the multi-task monitoring loop, several sources of uncertainty must be dealt with. These include human produced distortions in observing variables presented on a display, uncertainty in determining the status or state of the system, various hypotheses as to possible courses of action, and the difficulty in envisioning consequences of actions subject to the overall task objective. Clearly all of these sources of uncertainty are influential in determining the amount of workload imposed on the human.

In the present paper, a joint analytical - experimental sensitivity study was performed to investigate how changes in the information processing and decision-making load affects the decisionmaker's performance and operative behavior. The objectives of this effort were multifold:

- (1) To analyze the effect of variations in the workload imposed on the operator in regard to his performance and operational behavior.
- (2) To identify and understand limits on human information processing and action selection, especially in the case of very high workload, and attendant decreasing performance.
- (3) To try to find appropriate measures of a priori objective task difficulty and to compare them to previously known measures [4,7].
- (4) To validate the Dynamic Decision Model (DDM) predictions of human performance across a wider range of task parameters. Specifically, to do a model-data comparison while varying experimental parameters such as the number of task channels, the velocity of the tasks, the task values, and the processing time range.
- (5) To compare experimental results with subjective measures of workload. The technique used here is the SWAT (Subjective Workload Assessment Technique) [6].

The paper is organized as follows; in the remainder of this section a brief description of the Dynamic Decision Model (DDM) and the experimental paradigm used to validate it will be given. Part II will describe the simulations and experiments done for the sensitivity studies after which Part III will describe some of the main results in terms of model predictions and model data comparisons. In Part IV a study of a proposed a priori measure of task difficulty in terms of the ratio of average time required/time available will be presented. Part V discusses a comparison between experimental results and subjective workload prediction using SWAT. Comments on the implication of the results obtained on workload, human strategy changes and the DDM's predictions, as well as considerations leading to future research directions, will conclude the paper.

1.2 Review of the Dynamic Decision Model

Our main analytical tool will be the Dynamic Decision Model (DDM) previously developed by Patipatti, Kleinman and Ephraim [1-3]. This normative-descriptive model contains several interesting features. First, the analytic framework of DDM is based on optimal control, estimation, and semi-Markov decision process theories. Thus, this approach provides a general methodology for analyzing dynamic decisionmaking under uncertainty. Second the model introduces the important concepts called the "task" state and "decision" state. The task state is the detailed description of the internal variables associated with each task (position, velocity, type, etc.). The decision state, which is a functional transformation of the task state is the minimum number of variables that provides different information for making decisions (time available, time required to complete each task). Third, the DDM employs the widely - validated Kalman filter theory in an information processing submodel to provide the conditional mean and covariance of the decision state. Fourth, and perhaps most important, the DDM explicitly incorporates human limitations in the information processing as well as in the decision stages. These include the reaction time delays, randomness (i.e. observation noise), limited combinatorial capabilities and randomness in decisions.

A block diagram of the DDM is shown in Fig. 2. Each of the N tasks in the opportunity window is represented by a dynamic system acted on by disturbances to account for the nonstationarities in task characteristics. The human's perceived outputs $\{y_{p1}\}$ are delayed, noisy versions of the task states $\{x_{T1}\}$ and are contingent upon the monitoring process. The perceived outputs are processed to produce the best linear unbiased estimates of the task states $\{\hat{x}_{T1}\}$ and their associated covariances $\{E_1\}$ via a Kalman filter-predictor submodel previously used in the Optimal Control

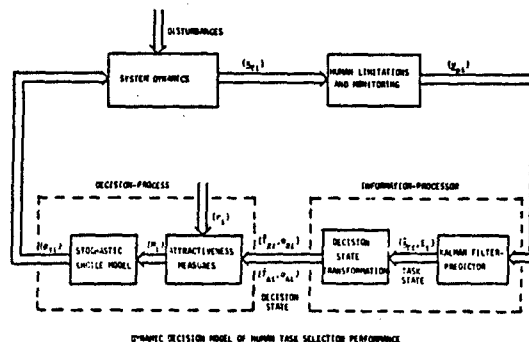


Fig. 2

Model (OCM) methodology. The statistics of the task states $\{\hat{x}_{T1}, E_1\}$ are, in turn, used to determine the first and second order statistics of the decision state $\{\hat{T}_{R1}, \sigma_{R1}\}$ time required and $\{\hat{T}_{A1}, \sigma_{A1}\}$ time available. The statistics of the decision states along with the task values, $r_1(t)$, are combined to determine the attractiveness measure, $M_1(t)$, of each task in the opportunity window. Subsequently, the measures are used to generate the probability $P_{d1}(t)$ of acting on each of the N tasks and the probability $P_{do}(t)$ of not acting on any task (or the monitoring probability, $P_{dm}(t)$).

The DDM is capable of predicting various performance measures, such as: the total reward earned, the percentage of total possible reward earned, the number of tasks processed, the total amount of time spent acting on tasks, etc. [1-3].

1.3 The Experimental Paradigm

A simple, yet realistic, computer controlled experimental set-up was considered as indicated in Fig.3.

In the experiments, the subjects observe a CRT screen on which multiple, concomitant tasks are represented by moving rectangular bars. The bars appear at the left edge of the screen and move at different velocities to the right, disappearing upon reaching the right edge. Thus, the screen width represents an "opportunity window".

The height of each bar corresponds to the reward (value) of the task. The amount of time required to process a task (in seconds) is represented by the number of dots displayed on a bar. A task is processed by the subject when he pushes the appropriate push-button as shown in Fig. 3. By processing a task successfully, the subject is credited with the corresponding reward (r_1), and the completed task is eliminated from the screen. An attempt by the decisionmaker to act on a task that cannot possibly be completed (i.e. the time required is greater than the time available) constitutes an error. No partial credit is given.

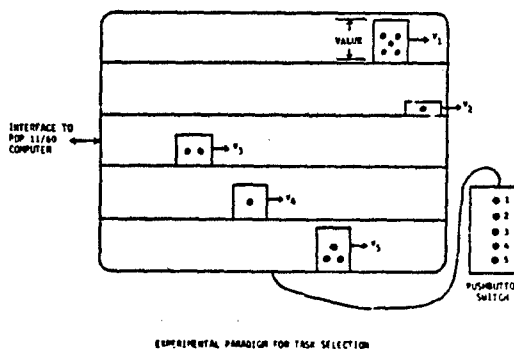


Fig. 3

In the present experimental paradigm, from which the DDM was developed, several factors influence the decisionmaker's workload. The human must process information involving the amount of time required (processing time dots), the value of a particular task (height), and the position and velocity of the task relative to the tasks's opportunity window deadline. All of these factors must be considered for each of the N task lines (channels) that are simultaneously vying for the decisionmaker's attention. By processing this information (i.e. reducing his uncertainty about the variables presented), the decisionmaker then uses his knowledge of the "state" of the system to develop an appropriate choice of which task to act upon.

II. METHODOLOGY FOR THE SENSITIVITY STUDY

II.1 Model Simulation

On the basis of existing notions concerning the effect of various task parameters on the human operator's performance and subsequent workload [4,7] four task attributes were considered. They include:

- A) different task velocities,
- B) different task values,
- C) different task processing times,
- D) varying the number of task lines (channels) to be monitored.

As indicated in Fig. 3, the task velocity is inversely related to the time available in which to successfully engage and process a task. More specifically, the time available to process task i at time t is given by

$$T_{a1}(t) = \frac{L - l_1(t)}{v_1(t)}$$

where

- L is the length of the opportunity window (1024 screen units) - 1 inch = 85 s.u.
- l_1 is the position of the task from the left edge (in screen units)
- v_1 is the velocity of the task (screen units/second)

It was postulated that by increasing the average task velocity a decrease in performance (as measured by the number of tasks completed and the total reward earned) would be obtained. A similar relationship between performance and the average processing time of the tasks was expected, since as more time was required to process a task fewer tasks could be completed prior to leaving the opportunity window.

By increasing the number of distinct task values or the number of task channels (lines) which must be monitored simultaneously, a decrease in performance was expected due to the increased demands placed on the information processing portion of the model in conjunction with increasing uncertainty (i.e. conflicts) among decision alternatives.

II.2 The Experimental Program

To parallel most of the sensitivity studies performed on the DDM by model simulation, the experimental paradigm described in I.3 was used. For each experimental setting, 3 different subjects, all graduate students in the Department of Electrical Engineering and Computer Science at the University of Connecticut, were tested on 2 different experiments (random arrivals) yielding 6 values for each average data point obtained. The duration of each experiment was 90 sec and the subject had immediate visual feedback on his overall performance. The subjects were trained on various experimental conditions; however data was taken only after a certain level of subjective confidence was felt by the subjects.

At the end of each sequence of experiments, a short

interview was designed to try to determine the various subjective strategies used. Overall, it was interesting to remark that although each of the subjects used different strategies to maximize their rewards, the intra and inter subject differences in performance were minor. Therefore, the experimental results presented in the next part will describe only the mean performance and not the inter subject variations.

III RESULTS AND MODEL - DATA COMPARISONS

Only the most significant results will be presented here. A complete presentation of the DDM sensitivity study can be found in [5].

III.1 The Sensitivity to Task Velocity

With regard to task velocity three separate DDM computer simulations were performed:

- 1) all tasks had the same velocity which was varied from 25 screen units (su) per second to 200 su/sec.
- 2) all tasks had the same mean velocity with uniform distribution range of +25 su/sec, the mean varied from 25 su/sec. to 175 su/sec.
- 3) all tasks had the same mean velocity of 100 su/sec with uniform distribution ranges of 0, +25, +50, +75, +100.

Results from part (1) are in general agreement with our hypothesis concerning the relationship between velocity and the time available to process a task. As expected a sharp decrease in the percentage reward earned was observed as the velocity of the tasks was increased. It does appear, however, that there is a tendency for the model to reach a "saturation point" as indicated by the number of tasks (20+2) completed and in the total expected reward earned during the simulation duration (90 sec).

In part (2) of the velocity study, the mean velocity was increased in a manner similar to part (1) but with some uncertainty in this value. Results were expected to show a general agreement with those obtained in part (1), which they did: The DDM performance was affected very little by the small uncertainty (+25%) on the task velocity.

Finally, part (3) attempted to investigate the effect of various amounts of uncertainty on the velocity of each task as perceived by the information processing portion of the model. A mean velocity of 100 su/sec was chosen for all the tasks with the interval over which the velocity was uniformly distributed (e.g. the velocity standard deviation) being increased. Results indicated that the model was able to successfully overcome the perceived uncertainty in task velocity and develop a strategy to improve performance, although this may have been a by-product of having more than enough time in which to process the tasks.

The experiments done for the sensitivity study on task velocity paralleled the simulation study (2) where all the tasks have the same mean velocity (with a small uncertainty of +25 su/sec) which was varied from "slow" to "fast". Here again the general hypothesis was that the performance of the subjects (in terms of percentage of reward earned) will degrade as the velocity of the tasks was increased, due to the corresponding decrease in the time available to process a task (i.e. the faster the task, the sooner it reaches the end of the opportunity window). Fig. 4 presents model/data comparison in terms of the percentage of reward earned and the total reward earned. Model predictions and experimental results were in excellent agreement up to the "nominal" situation of a mean velocity of 75 su/sec (a task of moderate velocity). Beyond that point, the human subject performed better than the model, reflecting a process of adjustment by the human to the increasing workload demands (an ability the model doesn't have at present).




Fig. 4 SENSITIVITY TO VELOCITY

In summarizing the velocity study, it seems apparent that the most influential aspect of velocity on performance manifests itself in the resulting time available/time required trade-off that must be evaluated in task selection. Clearly the information processing aspect of uncertainty in the perceived velocity played a lesser role in determining the resulting performance. This result is in agreement with existing notions regarding available time to process a task and the performance achieved, and forms a central theme of the workload theories which subscribe to the notion that the rate at which tasks must be processed is indicative of the workload imposed on the operator. In high workload situations, the divergence in performance between model predictions and human subjects seems to signal the fact that the human develops a strategy that adapts to high workload demands (high task speed, short inter-arrival times).

III.2 The Sensitivity to the Range of Task Values

In this study it was hypothesized that by increasing the number of possible reward values of the tasks a resulting increase in the uncertainty of selection due to conflicts would appear. Results of this part of the model simulation study showed consistent performance, as measured by the total percentage tasks completed, across the range of task values. It seems apparent that the model was able to successfully discriminate tasks of low value from those of high value and adopt a strategy reflecting this segregation.

As a way to verify the assumption that the model was able to discriminate high valued tasks from low valued tasks, a comparison of the actual reward earned and the mean value of the tasks multiplied by the number of tasks earned was performed (see Fig. 5). Clearly, here we see that the model was able to segregate the more desirable (high valued) tasks to be processed.

With regard to overall performance, a normalized reward (actual reward earned/average task value) was computed and plotted in (Fig. 6). Here again the consistent performance achieved by the model is exhibited.

Summarizing the study on the relationship between the performance achieved and the number of distinct task values, it appears that the model is able to suc-




Fig. 5 SENSITIVITY TO THE RANGE OF TASK VALUES

cessfully segregate more desirable tasks to be processed from lesser ones. By doing so, consistent performance is achieved. Hence it appears that the range of possible values produces a strategy which is able to compensate for any conflicts in choice such that a "saturation level" in performance is obtained.




Fig. 6 NORMALIZED REWARD

It appears that a heuristic measure of an a priori absolute attractiveness of a task i is a ratio γ_i

$$\gamma_i = \frac{\text{Value of task } i}{\text{Time required to process task } i}$$

which induces a multi-level threshold strategy:

1. for $\gamma_i > \gamma'$ "Always try to process task i "
2. for $\gamma_i < \gamma'$ "Never try to process task i "
3. for $\gamma_i \leq \gamma' < \gamma''$ "Depends on concurrent tasks"

For example: for the simulation with 10 different task values and 5 different processing times the values $\gamma' = 6.2$ and $\gamma'' = 4.3$ have been found.

Note: No experiment was performed in the study of sensitivity to the number of task values.

III.3 Sensitivity to the Range of Task Processing Time

In this study, the range of possible seconds needed to process a task was varied from all tasks having the same processing time of 1 second to that of the tasks having anywhere from 1 to 8 seconds processing time. The velocity of the tasks was chosen to be a nominal value (100 screen units/sec) corresponding to a moderately paced task with a small standard deviation about this value. Here again performance was expected to decrease as a result of the time available/time required trade-off that affects the decisionmaker's ability to process tasks prior to their deadlines. Indeed both model and experimental results confirm this hypothesis (Fig. 7).

As can be seen from the plots, the model and human are in good agreement up until the nominal case of tasks which had a processing time range of 1-5 seconds. After this nominal situation the human processed tasks




Fig. 7 SENSITIVITY TO PROCESSING TIME RANGE

with processing times which were below the average value of the tasks presented to him. This change in the type of tasks processed by the human reflects a change in strategy that accounts for the better performance achieved as compared with the model. In other words, the human adapts a more efficient behavior as the information load (range of possible processing times) increases. He discriminates among the tasks that are favorable (small processing time) from those deemed unfavorable (large processing time hence smaller time available to process the task prior to its deadline).

III. 4 Sensitivity to the Number of Task Channels

In this study, it was hypothesized that by increasing the number of task lines to be monitored simultaneously a decrease in performance would result due to increased demands on the information processing and action selection capabilities of the human. Fig. 8 shows the performance achieved by both the model and the human. As expected the percentage of reward earned decreases as more task lines are presented to the decisionmaker. The human again outperforms the model as the number of task lines increased beyond the nominal condition of five lines. This appears to be due to the fact that the human adjusts his behavior in an attempt to keep his performance constant when the number of tasks to be monitored simultaneously approached his short term memory capacity (5-7 lines).

The plot of the number of tasks completed by both the model and the human as a function of the number of task lines shows the so-called inverted-U shape hypothesized to exist under conditions of high workload. The experimental data, however, indicates a leveling off in the number of tasks completed by the human; achieved by adjusting his behavior to account for the increase in the number of task lines, (i.e. maintaining a subjective level of performance).

Also indicated in Fig. 8 is that the human is outperforming the model in terms of the probability of committing an error by the model is increasing as a function of the number of lines monitored, whereas the human maintains a relatively constant low error level. This fact can explain the constant human performance phenomenon exhibited earlier, in the sense that the

human, by trying to process less information as the workload increases, is able to commit fewer errors and therefore to achieve a higher performance than the model prediction.




Fig. 8 SENSITIVITY TO THE NUMBER OF CHANNELS

IV AN A PRIORI MEASURE OF TASK DIFFICULTY: THE ρ RATIO

As a result of the sensitivity program described before, in particular the study of sensitivities to velocity and processing time range, it was hypothesized that the ratio ρ

$$\rho = \frac{\bar{T}_R}{\bar{T}_A} = \frac{\text{Average time required}}{\text{Average time available}}$$

is a suitable measure of a priori task difficulty exhibited by a specific dynamic task in the current experimental paradigm. It can be easily shown that the ρ factor is given by the following formula:

$$\rho = \frac{n+1}{2n} \cdot \alpha(1-\beta^2)$$

where

- n is the maximum number of unit time (dots) per task
- $\alpha = \frac{\bar{V}}{V_{\max}}$ is normalized velocity
- \bar{V} is the average nominal velocity
- $V_{\max} = \frac{L}{n\delta}$ is a screen window limitation
- $\beta = \frac{\Delta V}{\bar{V}}$ represents the uncertainty on the velocity (uniform distribution)
- L : length of the opportunity window
- δ : time value of a unit time (in seconds)

Fig. 9 depicts a map of theoretical curves of equal ρ for different combinations of α and β . The hypothesis that the average factor $\rho = \bar{T}_R / \bar{T}_A$ correlates with workload and affects degradation of performance is verified by model simulation as well as by experimental results also shown in Fig. 9 for the case $n=5$, $\alpha=1/2$ (nominal). It can be seen that an increase in ρ leads

a decrease in performance.




Fig. 9 (a) ρ curves, (b) performance vs. ρ

A sensitivity study was performed on the model and also experimentally to check the assumption that for constant ρ the performance of the decisionmaker remains constant, although the range of dynamics parameters such as velocity and dot time value varied in a wide range. Fig. 10 confirms this hypothesis: the experimental data shows that the human develops a strategy to maintain his performance relatively constant when the average a priori workload (ρ factor) is kept constant. Note that when the value of the dots is $< .25$ sec, the performance degrades, probably due to the fact that we are very close to human time delay limitations. The model, however, does not achieve a constant performance level, but rather exhibits a degradation in performance. This degradation yields an increase in the average error probability unlike the relatively small and constant average error probability observed in the experimental data.




Fig. 10a

Fig. 10b

In conclusion the experimental results allow us to assume that the ρ factor can represent an appropriate a priori measure of task difficulty in the current dynamic environment.

V SUBJECTIVE WORKLOAD AND EFFECTIVENESS

In an effort to compare the level of workload imposed upon the human with that perceived by the human, rankings of subjective workload were calibrated with subject's rank ordering of workload along the three dimensions of Time Load, Mental Effort Load and Stress Load. The methodology used was the Subjective Workload Assessment Technique (SWAT) well documented in [6]. The experimental setting used was the paradigm presented in part I and the method of controlling imposed workload was by increasing the number of task channels from 2 to 10, all other parameters (velocities, task values and processing time ranges) remaining fixed. As expected, the average workload perceived by the subjects increased linearly with the number of task lines to be monitored (Fig. 11).





Fig. 11 PERCEIVED AND IMPOSED WORKLOAD

Corresponding to this increase in perceived workload, a change in human strategy under high workload conditions was noticed in conjunction with the performance level attained (recall Fig. 8). This change of strategy is depicted in Fig. 12, where an effectiveness factor ϵ given by

$$\epsilon = \left[\frac{\text{average value of tasks processed}}{\text{average value of tasks presented}} - 1 \right] \times 100$$

is plotted as a function of subjective workload for the same experiment. Essentially this effectiveness factor indicates that the human adapts a more discriminating strategy (acting on tasks of average task value higher than the average of those presented to him) when he is confronted with increasing information processing demands. This discriminating behavior accounts for the ability of the human to outperform the DDM as discussed earlier in adapting a strategy for high workload conditions. Fig. 12 shows clearly the 'jump' of operative behavior in terms of average effectiveness between low and high workload conditions.

Fig. 12 Effectiveness Factor vs. Subjective Workload



VI SUMMARY AND CONCLUSIONS

In this joint experimental-analytical study, performance obtained by the Dynamic Decision Model (DDM) and human subjects in the experimental paradigm of Fig. 3 was examined as a function of four task attributes which included different task velocities, different task values, different task processing times, and varying the number of task lines to be monitored. Results of the study showed a general agreement with existing hypotheses regarding human performance as a function of the four attributes mentioned. The primary factor affecting performance was that of task velocity as indicated in the time available/time required trade-off of traditional operator workload research. In addition, it was noted that the human operator was able to successfully discriminate various undesirable tasks from desirable ones as a means to keeping performance consistent. This may be due to a somewhat conservative threshold - strategy employed by the decisionmaker, but in general it is felt to be consistent with one's intuition.

The results of the model simulation study also indicate that a decrease in performance can be obtained not only from increased information processing demands, but also from situations in which action selection uncertainty (i.e. conflicts) exist. This was most clearly represented in the study involving the increased number of task lines to be monitored simultaneously. In that study it has been shown that beyond a certain number (~7) of task lines (channels), the operator performance decreases sharply due to the increasing conflict in action selection and limited information processing capability.

However, in the experiments, the subjects showed different performance in the high workload condition. Instead of decreasing, the subject's performance was maintained at a quasi-constant level due to adaptation in operative behavior in order to (1) process less information, (2) commit fewer errors, and (3) maintain a subjective level of performance across high workload environments. In diagram from the mechanism by which decisions are made by the human appears to be that of the following:




Fig. 13 Anatomy of a decision

Another contribution of the present effort is the proposition of an objective measure of a priori task difficulty in terms of the factor ρ . This measure ($\rho = T_R / t_A$) has been found to be adequate in the existing experimental paradigm, although more research will have to be done to improve the measure in different dynamic environments. Subjective workload was found to correlate very well with objective workload imposed by an experiment and an interesting find was that humans tend to be more efficient in reasonably high workload conditions.

REFERENCES

1. Pattipati, K.R.; Kleinman, D.L.; and Ephrath, A.R. "A Dynamic Decision Model of Human Task Selection Performance" *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. SMC-13 (2), March/April 1983.
2. Pattipati, K.R.; Kleinman, D.L.; and Ephrath, A.R. "A Dynamic Decision Model of Human Task Selection Performance", University of Connecticut Technical Report, Dept. of EECS, TR-82-1, March 1982.

3. Pattipati, K.R., "Dynamic Decision Making in Multi Task Environments: Theory and Experimental Results", Ph.D. Thesis, The University of Connecticut, 1980.
4. Moray, N. ed. *Mental Workload: Its Theory and Measurement*, New York, Plenum Press, 1979.
5. Soulsby, E.P.; Serfaty, D.; and Kleinman, D.L. "A Sensitivity Study of the Dynamic Decision Model", University of Connecticut Technical Report, Dept. of EECS, TR-83-12, July 1983.
6. Reid, G.B.; Shingledecker, C.A.; and Eggemeier, F.T. "Application of Conjoint Measurement to Workload Scale Development". In *Proc. of 1981 Human Factors Society Annual Meeting*, Oct. 1981 (a), pp. 522-526.
7. Soulsby, E.P. "Human Operator Workload: A Survey of Concepts and Assessment Techniques", University of Connecticut Technical Report, Dept. of EECS, TR-82-4, Nov. 1982.

INDIVIDUAL DIFFERENCES IN MILITARY DECISIONMAKING:
THE CLASSIFICATION PERFORMANCE OF ACTIVE SONAR OPERATORS

Joseph G. Wohl

ALPHATECH, Inc.
2 Burlington Executive Center
111 Middlesex Turnpike
Burlington, Massachusetts 01803

Introduction

The purpose of this paper is to demonstrate the fundamental importance of individual differences among military personnel in their information-seeking and processing behavior when faced with critical decisions in a realistic task. The data reported in this paper are collected at sea in the mid 1960s. The analysis of this data base has provided an object lesson showing the importance of empirical data to decisionmaking research.

Determinants of Decisionmaking Behavior

In general, the major determinants of decision-making behavior include such factors as information presentation rate, perceived problem complexity, perceived time available for decisionmaking, the number and quality of perceived alternatives, and the perceived risks (Sage [1], Sage and White [2], Pattipati, et al [3], and Wohl [5]).* In this paper, however, we shall examine two other factors for which empirical data are not often available. These are: 1) individual differences in cognitive and decision styles; and 2) level of expertise attained.

The Operational Problem in Active Sonar Classification

A modern active sonar, such as the SQS-23 or SQS-56, emits an extremely strong pulse of low frequency sound into the water on a nearly omni-directional basis. The outgoing sound pulse is diminished in intensity by both the inverse square law and the absorption of sound in water. It is reflected by surface wave fronts, bathypelagic fish, sea bottom anomalies, schools of shrimp, whales, and floating debris as well as real submarines, and is returned to the receiving equipment after having suffered another set of identical losses on the return trip. In the receiver, this signal information is mixed with noise from several sources including own ship and other ships as well as ocean sources.

The resulting mixture of signals and noise may be presented to the sonar operator in several display formats: a PPI display, a tactical range or time history recorder, and a pair of audio headphones. The sonar operator must examine the return information from each pulse or "ping", and as more and more information is accumulated on the sonar contact, he is required to make and report a decision as to whether the received information represents 1) a non-submarine contact, 2) a possible submarine contact, or 3) a probable submarine contact. This decision is made in conjunction with his watch supervisor and/or the antisubmarine warfare officer on board the ship.

This classification decision is only one of several activities which the sonar operator must perform. These activities include control of the sonar search process itself, monitoring of the display and detection of possible signals against the noise background, tracking of these signals over a period of time to determine their correlation and consistency, and ultimately classification itself.

The fundamental classification decision must answer the question, "Is this contact sufficiently like a submarine to be classified as a submarine?" There are five basic subquestions involved: 1) Is the contact truly moving? 2) What kind of reflective structure does it have? 3) How large is it? 4) What is its shape or aspect? 5) Does it have depth below the ocean surface? In addition, strong correlations among the answers to these five subquestions lend further credence to one or another of the possible hypotheses.

The Experiments

A U.S. Navy destroyer spent a number of weeks at sea during the mid-1960's collecting a representative sampling of submarine and nonsubmarine contact information on a modern active scanning sonar. Some sixty hours of contacts were recorded of which approximately 60% were submarine contacts and 40% nonsub contacts. The data were captured directly off of the sonar pre-amplifiers onto 54 reels of 56-channel tapes. The submarine contacts were made using friendly submarines under strict control in order to obtain information at various ranges, depths, speeds of both submarine and own ship, and submarine aspect angles, as well as under a variety and range of environmental and sea conditions. The nonsubmarine contacts included identified fish, porpoises, whales, bottom wrecks, and one floating log. The information was collated by a group at the Applied Research Laboratory of the University of Texas.

From this collated information a realistic, 100-item test was constructed consisting of fifty subs and 50 nonsub contacts and as much as 25 to 40 "pings" per contact, all taken under a representative range of sea conditions including sea state, wind velocity, wind direction, layer depth, bottom type, and bottom depth. The test tapes were then presented under highly controlled conditions on real sonar equipment to 37 Pacific Fleet sonar operators. These subjects represented a wide range of experience, measured in terms of years of Naval service, number of sonar teaching jobs, number of sonar schools attended, number of ship assignments, total years of sonar experience and years of experience with the sonar used in this experiment.

In addition to the regular sonar equipment with its displays and controls, the subjects were given ten push buttons labeled 0 through 4 and 6 through 10 and instructed to press number 10 if absolutely certain of a submarine, number 0 if absolutely certain of a

*References are indicated by number in square brackets and appear at the end of this paper.

nonsubmarine, and numbers 1 through 9 as appropriate on each ping. Omitting the number 5 from the group of push buttons essentially produced a scaled forced choice situation. Note that pushing a button is essentially the same as a subject reporting his posterior probability that his response resulted from the presence of a submarine. Thus, this method is a mechanization of Green and Swets' [5] subjective confidence rating method for deriving Receiver Operating Characteristic (ROC) curves. It effectively compresses multiple unknown dimensions (i.e., the various sonar cues and the weights given to them by the subjects) into a single subjective dimension. Both the experiment and the results are described in Wohl et al [6] and Nacht and Wohl [7].

Experimental Results

Using this method, ROC curves were developed for each of the 37 sonar operators. It was quickly apparent that the operators fell into two groups which could be labeled "average" and "expert". The single exception to this categorization was the one subject who exhibited a negative ROC curve. Figure 1 shows the combined results for the "average" group of 33 operators. Note that the performance is little better than random, and that accumulation of additional information had no effect on performance for this group.




Figure 1. Average ROC Curves of "Average" Group

$P(s/S)$ = probability that sonar operator says "sub" when contact is a sub ("hit" probability)

$P(s/NS)$ = probability that sonar operator says "sub" when contact is a non-sub ("false alarm" probability)

By way of contrast, the average curves of the highest performance or "expert" group of 3 operators is shown in figure 2. Here it is clear that these operators were able to make use of the additional information in improving their discrimination capability.

In order to determine the effect of individual differences on classification performance, the data for each subject were aggregated in the following way. Pressing any of the buttons 0 through 4 was taken as a nonsub response, while pressing any of the buttons 6




Figure 2. Average ROC Curves of Highest Performance or "Expert" Group

$P(s/S)$ = probability that sonar operator says "sub" when contact is a sub ("hit" probability)

$P(s/NS)$ = probability that sonar operator says "sub" when contact is a non-sub ("false alarm" probability)

through 10 was taken as a sub response. The performance results were then plotted in terms of "hit" versus "false alarm" scores. These two dependent variable scores were then subjected to multiple correlation analysis with the six independent variables of training and experience mentioned previously. Unfortunately, training and experience data for seven of the 37 subjects were not available. For the remaining 30, the analysis results showed an absence of significant correlation between performance and experience.

The data for the 30 remaining subjects were then divided into two groups labeled "unbiased" and "biased". The criterion for categorization is evident in the differences between figures 3 and 4. Figure 3 shows the aggregated classification performance for a typical "unbiased" sonar operator. The point labeled number 1 in figure 3 represents the average performance for the first "ping" across all 100 test items for that subject, point number 2 for the second "ping", and so forth. Of the 30 subjects for whom experience data were available, 10 showed "unbiased" characteristics similar to figure 3 while 20 showed "biased" characteristics similar to figure 4. The term "biased" is used here to indicate the fact that operators exhibiting this characteristic almost always reported "nonsub" on the first "ping"; i.e., they would almost always tend to push buttons 0 through 4 on the first "ping". In addition, they exhibited a delay of between 5 and 15 "pings" before their bias disappeared; i.e., they seemed to require an accumulation of between 5 and 15 "pings" worth of sonar information to overcome their bias.

At first this was believed to be the normal result of more extensive experience at sea on the part of the "biased" operators. Since most at-sea sonar contacts with active sonars are indeed found to be nonsubs, one would expect the development of such a base rate bias over a period of long service.




Figure 3. Classification Performance Curve of Typical "Unbiased" Sonar Operator




Figure 5. Correlation of Experience with Classification Performance For "Biased" and "Unbiased" Sonar Operators




Figure 4. Classification Performance Curve of Typical "Biased" Sonar Operator

This hypothesis was then tested by subdividing the group of 30 sonar operators into "biased" and "unbiased" subgroups and again running the multiple correlation analysis with the six experience factors as before. The results of this analysis are shown in figure 5. Here, the multiple correlation coefficient is plotted against the accumulation of information in terms of number of "pings" on each contact. From figure 5, it is clear that the "unbiased" subjects (represented by the upper curve scores) showed high and significant correlation with experience, regardless of information accumulation, as might be expected of "normal" people. On the other hand, the 20 "biased" subjects showed a low correlation with experience until at least 5 to 15 "pings" of information had been accumulated.

Discussion and Conclusions

From the foregoing data analysis, several conclusions can be reached. First, "bias" is not correlated with experience; rather it appears to be inherent in the subject. Secondly, the presence of bias results in minutes of lost time in an operational setting; it should be noted that 5 to 15 "pings" represents an additional 3000 to 9000 yards headway made by an attacking submarine before action is taken. Thirdly, even though "biased" operators are evidently trying to reduce their false alarm rate (see figure 4), the correlation results shown in figure 5 indicate that biased operators were just as biased about hits as about false alarms.

It is clear that a deeper understanding of this type of judgment bias and its operational impact is critical to a number of naval activities including personnel selection and assignment, command and control system design, and decision support system design. Since passive sonar is the primary ocean surveillance in current use today, the study results reported herein are primarily of historical interest. But they demonstrate the critical role that empirical data can and must play in understanding human-machine interaction and in designing for it. Such data are needed for today's critical Navy decisionmaking tasks such as those of the specialized Warfare Commanders, the Combined Warfare Commander, the Tactical Action Officer, and those who select and/or modify firing doctrine and rules of engagement. Such data will help us not only to better understand these decisionmaking functions, but will also provide an improved basis for 1) building and validating decision models, 2) designing better systems, and 3) making improved human and system employment decisions.

References

1. Sage, A.P., "Behavioral and Organizational Considerations in the Design of Information Systems and Processes for Planning and Decision Support," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-11, No. 9, Sept. 1981, pp. 640-678.
2. Sage, A.P. and E.B. White, "Decision and Information Structures in Regret Models of Judgment and Choice," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-13, No. 2, March/April 1983, pp. 136-144.
3. Pattipati, K.R., D.L. Kleinman and A.E. Ephrath, "A Dynamic Decision Model of Human Task Selection Performance," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-13, No. 2, March/April 1983, pp. 145-166.
4. Wohl, J.G., "Force Management Decision Requirements for Air Force Tactical Command and Control," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-11, No. 9, Sept. 1981, pp. 618-639.
5. Green, D.M. and J.A. Swets, Signal Detection Theory and Psychophysics, Wiley Press, New York, 1966.
6. Wohl, J.G., D. Promisel, S. Barnes, L. Stowens, M. Nacht and H. Henderson, "Active Sonar Classification Using the AN/SQS-23 Sonar," Final Report No. BSD 69-707, Dunlap and Associates, Inc., Darien, Conn., 16 May 1969, DDC #AD503129 (Confidential).
7. Nacht, M., and J.G. Wohl, "Appendix II: Data Analysis From an Experiment Conducted by Applied Research Laboratory of the University of Texas at Austin" (Unclassified Appendix to Reference [6] above).

Dup^e

THE COGNITIVE ORGANIZATION OF SUBMARINE SONAR INFORMATION:
A MULTIDIMENSIONAL SCALING ANALYSIS

Kevin Laxar, George Moeller, and William H. Rogers

Behavioral Sciences Department
Naval Submarine Medical Research Laboratory
Box 900 SUBASE NLON, Groton, CT 06349

Summary -- Nonmetric multidimensional scaling (MDS) techniques were employed to determine how sonar information is organized and assigned priorities by the Submarine Conning Officer (CONN). Data were collected from 95 Submarine Officers with varying amounts of at-sea experience. All types of information proposed for display in modern sonar systems were classified by the investigators into 15 categories. Descriptions of the categories comprised the stimuli for the two tasks the subjects performed. In an unconstrained sorting task subjects sorted the 15 stimuli into groups according to similarity of the sonar information described, to provide data for the MDS analysis. In a ranking task, subjects rank ordered the stimuli according to importance at CONN. The MDS analysis provided evidence that the officers organize sonar information in two dimensions, related to Information Source and Information Destination, while the rank order data indicated that most importance was attached to the information at the extremes of these dimensions. Significant agreement was found among all subjects, regardless of experience level, in the way the sonar information was psychologically organized and prioritized.

Recent technological developments have made the human-machine interface increasingly complex in terms of the kind and amount of information available and how it is displayed. The Naval Submarine Medical Research Laboratory has addressed some of the resultant problems in the design and operation of automated information systems. In particular, submarine sonar systems. One phase of this project has been to identify those pieces of sonar information that are perceived by the Conning Officer (at the "CONN"), who is immediately in charge of ship operations, to be most useful in ship control. Current hardware makes it possible to display any or all sonar information at the CONN, from raw auditory data to refined visual displays of predicted ships' positions. Two important considerations, however, may dictate that less information be provided than is technologically possible. One of these is financial, in terms of hardware and software costs. The other, which this research addresses, is the limitation by human information processing capacities, since many information processing theorists consider too much information a source of performance degradation (e.g., [1]).

There exist a number of different approaches to identifying and prioritizing the sonar information that should be displayed at CONN, but each has its associated problems. It has been our experience, for example, that judgments by systems engineers frequently are not well received by the operational forces, and that simple polls of experienced submarine officers often yield equivocal results. More meaningful information, on the other hand, could be obtained by empirical assessment of alternatives during real or simulated operations, but such an approach can be expensive and time-consuming. In lieu of these approaches, Zachary [2] has employed nonmetric multi-

dimensional scaling techniques [3]-[6] in the context of Naval Air antisubmarine warfare (ASW) in prioritizing decision-making situations. Such techniques have been applied in the present study to judgments about sonar information, to determine how such information is organized and assigned priorities by the submarine Conning Officer.

Method

Subjects

Data were collected from 95 Naval Officers in the New London area. In order of decreasing seniority and experience, the sample consisted of 11 Commanding Officers, 16 Executive Officers, and 30 junior men qualified as Officers of the Deck, from eight fast attack (SSN) and eight fleet ballistic missile (FBM) submarines. In addition, 38 junior officers, who had recently completed the Submarine Officer's Basic Course at the Naval Submarine School, participated. This last group, in general, had no at-sea experience.

Stimuli

The various types of information available from current and proposed sonar systems were classified by the investigators into 15 categories, as listed in Table I. Descriptions of these categories comprised the 15 stimuli for the tasks to be performed. Discussions with sonar instructors indicated that the selected categories were exhaustive of the types of sonar information that could be presented at CONN. Each of the stimuli was typed onto a separate card, numbered on the reverse side, to create the stimulus deck. A questionnaire administered after the data collection confirmed that the categories were meaningful and that no important piece of information had been omitted.

Procedure

To provide data for the multidimensional scaling analysis, subjects were first asked to perform an unconstrained sorting task, arranging the stimuli into as many or as few groups as they felt necessary, according to similarity. The definition of similarity was left up to the subject. Cards which described similar categories were to be placed in the same group, and any card which described a unique category was to be placed by itself. Then, to provide additional data for interpreting the scaling analyses, subjects were asked to rank order the stimuli according to importance at the CONN for two different operational missions. The first mission assumed an SSN on an ASW direct support patrol. In such a scenario, own ship would seek out and follow enemy submarines. The second mission assumed an FBM patrol in an area where a high density of sonar contacts was expected. In this scenario, own ship would remain in a designated area and try to avoid detection by enemy vessels. Subjects were instructed

then to mark their rank-ordered lists to show which options were necessary, merely desirable, or unnecessary. After each task, subjects recorded their data on answer sheets according to the code number on the back of each stimulus card.

Due to time constraints in obtaining data from these subjects, the sorting procedure was used in lieu of the pairwise judgment of similarity often employed in this type of analysis. Data were usually collected from small groups of subjects, such as one ship's crew, in sessions lasting approximately one hour.

Results and Discussion

The data from all subjects for the unconstrained sorting task were entered into a computer program which produced a dissimilarities matrix for the 15 stimuli, assigning values to the 105 pairs of stimuli according to the number of times subjects placed them in the same group. This initial procedure thus produced a proximities matrix from the nominal scale sorting data. The resultant matrix, in turn, was the input to the KYST-2A multidimensional scaling program [7]. Through this technique as employed here, a configuration of points (stimuli) in Euclidean space is constructed by an iterative adjustment process, based on the observed dissimilarity between all pairs of stimuli. The final configuration is then rotated so that the principal components of the points lie along the coordinate axes. The object of this procedure is to help determine the underlying psychological structure of the stimulus domain, namely the various pieces of sonar information.

TABLE I

The categories of sonar information that comprised the stimuli. The numbers indicate the aggregate rank ordering by importance, and the partitions indicate the degree of necessity, for all subjects.

<u>Rank</u>	<u>Kind of Information</u>	<u>Necessity</u>
1.	Contact Summary - Geographic	NECESSARY
2.	Contact Summary - Tabular	
3.	Single Contact Data	
4.	Tactical Aids	
5.	Own Ship Data	

6.	Contact's Active Sonar	DESIRABLE
7.	Raw Visual Displays	
8.	Ocean Acoustic Parameters	
9.	Ranging Data	
10.	Classification Aids	
11.	Environmental Parameters	

12.	Passive Sonar Setup	UNNECESSARY
13.	Raw Auditory Signals	
14.	Active Sonar Setup	
15.	Sonar Hardware Status	

The computer analysis was repeated with 10 different starting configurations to ensure that the obtained solution was a result of the stress value reaching a global, rather than a local, minimum.

The resulting two-dimensional solution is presented in Figure 1, with the number of dimensions selected according to the suggestions given by Shepard [8]. These included consideration of data values in the dissimilarities matrix, stress values for other dimensionalities, and meaningfulness in the interpretation of the axes.

The labeling of the dimensions in a multidimensional scaling configuration is, for the most part, based on the information available to the investigators about the set of stimuli being scaled. Examination of Figure 1 leads us to believe that, at least for the sorting data we obtained, the officers organize sonar information in terms of no more than two basic dimensions: data concerning sources of information, as shown along the vertical axis, and data related to the destination of that information, as given along the horizontal axis. The two extremes of the Information Source dimension are delimited by information from the world external to the submarine. At one end are auditory and visual displays of the relatively unprocessed sonar signals arriving at the ship's hydrophone arrays, obtained in passive mode from noise generated by the sonar contact. Also here lies information about the contact derived from any active sonar transmission the contact makes. At the other end of the scale is information about the environment which bears on sonar performance, such as sea state, ocean depth, and computed parameters for the acoustic properties of the surrounding ocean area. Information about, and derived by, own ship lies between the outside world extrema. Hence, this axis can be labeled as Contact versus Environment.

The Information Destination axis is concerned with where in own ship, the submariner's inside world, the available information is directed. The axis is delimited at one end by factors relevant to the CONN, which influence the maneuvering of own ship: a table listing all contacts and their classification, such as friendly or hostile, surface or submerged; a geographic picture showing the positions of contacts in relation to own ship; and displays showing predicted future positions of contacts and the effects of trial maneuvers. At the other end of the dimension is information relevant to the sonar personnel: the status of own ship's sonar equipment (performance monitoring/fault location) and the current utilization of the various pieces of active and passive sonar equipment. This axis can therefore be labeled in terms of Sonar versus CONN.

The data of all four groups of subjects were aggregated for each ranking task according to importance, and the Kendall coefficients of concordance W [9] were computed to assess between-judge agreement. For both rankings, agreement was highly significant, as indicated by the chi-square test. For the situation in which the submarine was acting in an ASW support role, a coefficient of $W = .46$ was obtained, $X^2(14, N = 94) = 609.9, p < .001$. For the FBM patrol situation, a coefficient of $W = .44$ was obtained, $X^2(14, N = 94) = 577.9, p < .001$.

The rank order of importance for the 15 items of sonar information was determined from the sum of their ranks from all subjects. This ranking for the FBM patrol situation is given in Table I, with the partitions according to necessity indicated. The ranking for the ASW support role situation was identical except for a transposition of items ranked 13th and 14th, and

when both rankings were combined, the ranking was as shown in the table. The ordering and categorization of these stimuli according to their importance at CONN appears quite reasonable. Those items deemed necessary are exactly those important to maneuvering own ship: the location of sonar contacts in relation to own ship, the classification of each contact, and the motion of own ship. Those items ranked moderate in importance were described as desirable, or nice to have, but not absolutely necessary. These items appear to be ones which are less useful, in themselves, to CONN in operating own ship, but which may help evaluate the quality of information categorized as Necessary. In that regard, it is quite unexpected to find the visual displays of relatively unprocessed sonar signals to be ranked as high as seventh. This may indicate a tendency of CONN to "look over the shoulder" of those in Sonar, perhaps just to make sure Sonar is not missing any contacts. Finally, those items labeled as Unnecessary are those concerned with the operation of the sonar system, generally under the complete purview of the Sonar Supervisor.

The numbers in Figure 1 show this rank order written beside the labeled points on the two-dimensional scaling configuration. Those four items ranked most important to, and directly concerned with, the function of CONN, are located together at the appropriate end of the Information Destination dimension. These are followed by data about Own Ship and Contact's Active Sonar, both slightly removed from the CONN extremum and relatively distant from each other, in the directions of the ends of the Information Source dimension. The next two items in importance are very close to the extremes of the Information Source dimension, the Raw Visual Displays at the Contact end, and Ocean Acoustic Parameters at the Environment end. Those items ranked least important relate to the sonar equipment and are placed at the appropriate end of the Information Destination dimension.

It may be noted, however, that the ranking by importance follows, in some approximate manner, the arrangement of the stimuli as one proceeds along the Information Destination dimension from CONN to Sonar. To determine if this unidimensional ranking formed the underlying basis for the configuration given by the KYST-2A scaling, the program was run again using the rank order as the starting configuration for a one-dimensional solution. As with other hypothesized unidimensional starting configurations run previously, the stress value for the one-dimensional solution was not improved beyond the value originally obtained. This result further indicates that while a meaningful unidimensional ordering can be imposed on these stimuli, the underlying organization is yet two-dimensional. In addition, however, Information Destination is very likely the more salient of the two dimensions.

To determine if the four groups of officers had organized or ranked the sonar information differently, complete separate analyses as described above for all subjects were computed on the data from each group. In all cases, results indicate that a two-dimensional solution was most appropriate. The KYST-2A scaling configurations were very similar for all groups, with the stimulus points in slightly different positions in their respective quadrants from one group to another. The one exception was that the Executive Officers placed Own Ship Data closer to the Contact rather than the Environment end of the Information Source dimension.

As indicated by the significant coefficient of concordance given above for all subjects, rankings between groups were also rather similar, characterized for the most part by transpositions of adjacent stimuli from one group to another. The notable exceptions were that the Commanding Officers ranked the Raw Visual Displays second, in the Necessary category, and the recent graduates of Submarine School ranked that same information in the 12th position, in the Unnecessary category. Commanding Officers were perhaps reflecting the desire to monitor the raw data in order to confirm

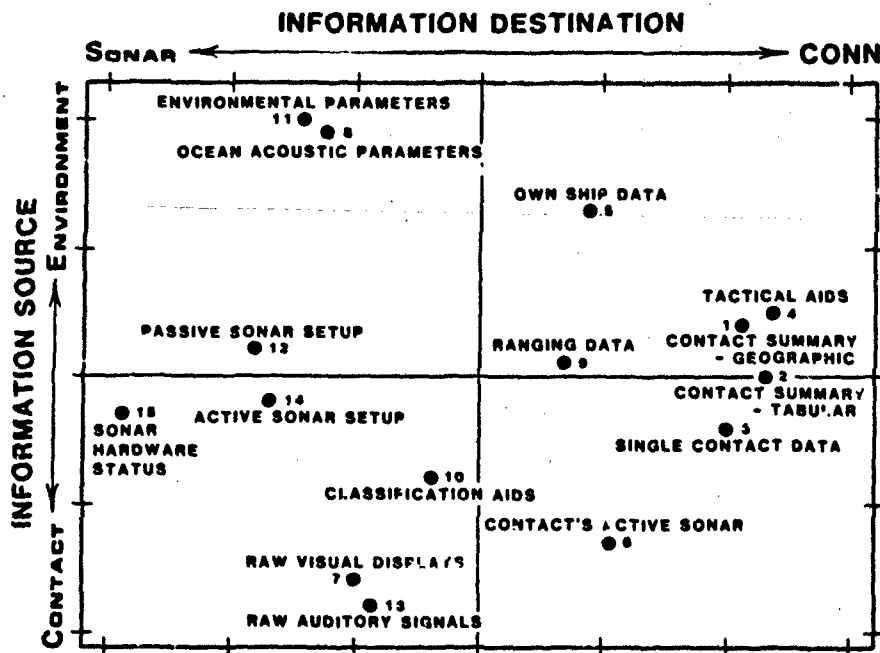


Figure 1. The two-dimensional solution for the KYST-2A scaling analysis, for all subjects. The numbers indicate the ranking by importance for the 15 categories of sonar information.

the accuracy of inferences represented by the categories, or to be closely involved with all phases of the ship's operation. The least experienced group, on the other hand, may have been expressing recently acquired training doctrine. Between-judge agreement for all groups was highly significant, with coefficients of concordance of $W = .42$ to $.59$ obtained. Within groups, as well, little difference was seen in rankings for the two different tactical missions.

When the rankings were compared between officers assigned to FBM submarines and those assigned to SSN ships, again, little difference was evident in the way the two groups ranked the stimuli for the two operational missions, and the rankings followed the same general pattern as presented in Table I. In addition to minor reversals in rankings between the two groups of officers, however, the FBM officers consistently ranked Own Ship Data as more important, third over all, than did the SSN officers, who ranked it seventh. This difference may reflect greater general concern on the part of the FBM officers with their ship being "on station," consistent with the mission of an FBM patrol. Similarly, SSN officers ranked Single Contact Data and Contact's Active Sonar two positions more important than the FBM officers did, possibly reflecting consistency with the SSN's mission. It should be noted that the differences between these groups are minimized by the fact that the officers could have had a varied range of experience on a submarine other than the type to which they were currently assigned.

The rankings for the groups of FBM and SSN submarine officers were combined with each other and across the two types of missions, as well. When compared with the ranking from the Submarine School graduates, the latter attached more importance to the Ocean Acoustic and Environmental Parameters and less importance to Own Ship Data and, as noted above, Raw Visual Displays. It is suggested that these differences may reflect experience gained at sea versus the aspects of ship operations emphasized in the Submarine Officer's Basic Course.

Summary and Conclusions

This study represents a successful application of the multidimensional scaling model, providing a representation of the way in which various pieces of sonar information are organized in the mind of the submarine Conning Officer. Results indicate that there is substantial agreement among officers of various levels of experience regarding the way the kinds of sonar information are organized. There is also agreement among these groups in the relative importance of these pieces of information in two different operational scenarios, both of which yielded similar rankings.

At least for the data obtained from unconstrained sorting by similarity, multidimensional scaling analyses suggest that two dimensions, at most, are required to describe the Conning Officers' conceptualization of the relations among various types of sonar information. One dimension is related to the source of available sonar information, whereas the orthogonal, and primary, dimension relates to where in own ship that information is directed or handled. The former dimension is laid out according to information from the sonar contact, from own ship, and from the ocean environment. The primary dimension involves sonar operations at one end and Conning Officer's responsibilities at the other.

When ranked according to importance, the information that the officers appear to require most is that from the extremes of the dimensional axes, except for

that information directly concerned with the sonar hardware. For system design, these results suggest that data about the sonar system are least desired at CONN and hence could be omitted from the CONN's display if financial or information processing limitations dictate that all information should not be made available. If further restriction of kinds of data to be displayed at CONN were necessary, investigation of those types of information in closest proximity to each other in the multidimensional scaling solution, indicating highly similar data, would be appropriate to determine if there are any completely redundant displays. An hierarchical clustering analysis is underway to assess this redundancy.

Acknowledgments

The authors wish to express their gratitude to Dr. Thomas Santoro for his computer assistance, and to Dr. Phipps Arable for his most helpful consultations on various aspects of the multidimensional scaling analyses.

This research was performed under the Naval Medical Research and Development Command Research Work Unit M0100-PN-001-1015. The opinions and assertions contained herein are those of the authors and should not be construed as official or necessarily reflecting those of the Department of the Navy or the Naval Submarine Medical Research Laboratory.

References

- [1] I. L. Janis and L. Mann, Decision making. New York: Free Press, 1977, p. 16.
- [2] W. W. Zachary, Application of multidimensional scaling to decision situation prioritization and decision aid design. Technical Report 1366-B. Willow Grove, PA: Analytics, 1980.
- [3] J. B. Kruskal, "Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis," Psychometrika, vol. 29, pp. 1-27, 1964.
- [4] J. B. Kruskal, "Nonmetric multidimensional scaling: A numerical method," Psychometrika, vol. 29, pp. 115-129, 1964.
- [5] R. N. Shepard, "Analysis of proximities: Multidimensional scaling with an unknown distance function. I.," Psychometrika, vol. 27, pp. 125-140, 1962.
- [6] R. N. Shepard, "Analysis of proximities: Multidimensional scaling with an unknown distance function. II.," Psychometrika, vol. 27, pp. 219-246, 1962.
- [7] J. B. Kruskal, F. W. Young, and J. B. Seery, How to use KYST-2, a very flexible program to do multidimensional scaling and unfolding. Murray Hill, NJ: Bell Telephone Laboratories, 1977.
- [8] R. N. Shepard, "Representation of structure in similarity data: Problems and prospects," Psychometrika, vol. 39, pp. 373-421, 1974.
- [9] M. G. Kendall, Rank correlation methods. London: Griffin, 1948, chap. 6.

DECISION AND DISPLAY ANALYSIS IN A SIMPLE SURVEILLANCE PROBLEM

Frank L. Greitzer and Ramon L. Hershman

Navy Personnel Research and Development Center
San Diego, CA 92152

Summary

Last year at this workshop we reported on human performance in a decision making task posed in terms of target surveillance. The target was either at a previous fix or had moved to a new location offset by a known distance and random angle. The problem was to decide, based on a noisy sample of data, which of these two states was true. To assess human information processing abilities we derived an optimal processor to compare with human performance data. In this report we pursue the relationship between visual and auditory representations of the task and describe effects of alternative representations on human performance.

Problem Definition

Details of the problem and its analysis were described earlier [1]. Briefly, we assume that the target is stationary when the data are observed and that its position is either fixed at an arbitrarily defined origin or exactly R units removed in any direction. Let N observations be presented as dots on a CRT display, where each dot encodes the reported (x,y) coordinates. Let sensor error be produced by the circular normal density. Then if the target hasn't moved, the dots will tend to cluster around the origin. If it has moved, the dots tend to cluster around its new locus $(R\cos\theta, R\sin\theta)$ where R is the fixed offset distance and θ is taken to be a uniform variable on $(0,2\pi)$.

Thus, the decision maker must decide between only two possible states of nature, S_0 and S_R . In state S_0 the target is at $(0,0)$; in state S_R it has moved R units away. Figure 1 (a) shows some typical stimuli in the original "DOTS" representation for samples of size 7 from S_R (signal) and S_0 (noise), respectively. On each trial the observer sees one such stimulus or the other and must decide whether S_0 or S_R is true. Stimuli b-e in Figure 1 are other visual representations of the identical data.

Optimal Performance

Requiring the decision maker to minimize the expected cost of decisions, we derived the optimal Bayes procedure and the optimum rule (see [1] for details). viz.:

Decide S_R if and only if $d > d^*$,

where d is the observed distance from the origin to the center of mass (centroid) of the data; and the criterion d^* contains a modified Bessel function of order zero and depends on the costs, priors, sample size N , error variance σ^2 , and R . The criterion d^* is a decreasing function of N/σ^2 and, for symmetric costs and priors, approaches $R/2$ in the limit.




Figure 1. Visual displays used in the decision making experiments. The five stimuli shown as "signals" (and those as "noise") represent identical samples of size 7. All of the stimuli encode the same information; viz., in polar coordinates the encoding is (r_i, θ_i) representing range and bearing from the origin.

Explanation of stimuli: (a) DOTS--original spatial representation of points in the plane. (b) SPIKES--spikes are drawn at bearings θ_i and with lengths proportional to r_i . Size of circle is arbitrary. (c) GLYPHS--same as SPIKES except rays are used. (d) WEDGES--wedges extend from the origin to each sampled point. (e) ZAGS--vectors of length r_i and orientation θ_i are successively appended.

In the context of the theory of signal detectability [2] we regard the movement of the target as a "signal" to be detected and obtain the probabilities of detection and false alarm. These are, respectively, the probabilities that the observed distance d exceeds d^* , given S_R and S_0 . Figure 2 shows theoretical Receiver Operating Characteristic (ROC) curves for sample sizes $N = 1, 3, 5, 7$, and 9 with $R = 1.2$ units and $\sigma^2 = 1$. For a given sample size, a decision maker who uses the distance statistic d is constrained to the given curve.

Human Performance

Spatial Representations

We assigned equal priors and collected performance data for the DOTS task with sample sizes $N = 1, 3, 5, 7$, and 9 ; $\sigma^2 = 1$. A microcomputer displayed, on each trial, the circle of radius $R = 1.2$ units (2.2 cm) and plotted the sampled data as points in the plane. We used both yes/no and rating-scale procedures (see [2], pp. 32-43) for data collection. Feedback--right or wrong--and the correct location of the target were given after each play. Three observers were tested for approximately 500 plays at each sample size; the last 200 plays were analyzed. Representative rating-scale ROC curves for one observer are plotted as dashed lines in Figure 2. We note that our observers only approximated the optimal procedure, and although they improved as N increased, they failed to extract all the available information.

We also tested other visual representations of the data. In this regard, any information display constitutes some level of decision aiding. At one extreme would be a display of numbers for the x, y coordinates. At the other extreme would be simply showing the optimal decision. Our displays in Figure 1 lie between these. We sought alternative representations that preserved the dimensions of the original task. Thus, every display codes distance and angle information for each observation--these dimensions remain unintegrated and their processing is left to the observer. Mathematically, the task is identical in every case. However, the cognitive processing required seems quite different: Shapes of stimuli $b-e$ seem more salient than distances to centroids.

We collected performance data for the five displays in Figure 1 and found little, if any, differences; the ROC curves appear in Figure 3. The failure to find differences, we believe, attests to the versatility of human information processing. We note, however, that observers had unlimited time, no stress, and no additional workload conditions--perhaps such manipulations would yield differences in performance.

Acoustic Representation

As we noted last year, the solution to our spatial task is identical to the solution of a well-known problem in acoustic and radar signal processing: detecting in Gaussian noise a sinusoid signal known exactly except




Figure 2. Optimal (solid curves) and human (open circles) ROC curves for sample sizes $N = 1, 3, 5, 7$, and 9 ; observed curves are each based on 200 trials.




Figure 3. Observed ROC curves for the five information displays with $N = 5$; $R = 1.2$, $\sigma^2 = 1$. Each curve is based on 200 trials with a rating-scale procedure. [D = DOTS, S = SPIKES, G = GLYPHS, W = WEDGES, Z = ZAGS]

for phase. Here the input is passed through a narrow-band filter tuned to the signal frequency; the amplitude of the resultant envelope is then submitted to a criterion device to reach a decision. There is a strict isomorphism between this acoustic problem and the spatial task for our observers. The unknown phase of the sinusoid is precisely the unknown angle θ of the target's movement in the spatial domain. The amplitude of the envelope output by the narrow-band filter is the distance statistic d in our task. We find no prior reference to this striking isomorphism in either the engineering or behavioral science literature. We were particularly interested in comparing human performance with different sensory modalities on these mathematically identical tasks.

A psychoacoustic experiment reported in 1964 [3] (see also [2] and [4]) met the conditions of our task and essentially played the spatial problem through the observer's ears. A sliding rating scale was used to report the observer's confidence that the sine wave was present and this afforded a precise plot of the ROC curve. One of the conditions in [3] had acoustic parameters $E_s/N_0 = 14$ that are mathematically equivalent to a sample size of 19.4 in our spatial task [5]. The auditory ROC curve is shown in Figure 4 (open circles). Note that the acoustic task was much easier (in theory) than any of our conditions. However, the auditory ROC curve approximates the theoretical curve that we derived for our task with $N = 3$; it also is close to our observed data for $N = 3$. These comparisons indicate a severe loss in sensitivity in the acoustic modality.

It is clear that humans perform far better in the spatial equivalent of the detection problem. We attribute this superiority in part to the human's limited ability to process acoustic phase information [6]. Coherence in phase, of course, speaks to the presence of the signal--and in the spatial domain the variation of phase angle for the N observations is readily observed. However, this information is not accessible acoustically





Figure 4. Observed ROC curves for the auditory detection experiment (open circles, after [3]); and the isomorphic visual task for $N = 20$ with no phase information (filled circles).

because the human auditory system is, we might say, phase-deaf: at least with regard to differentiating simple pure tones differing only in phase. So, we next considered withholding phase information in the visual task.

Comparing Modalities

We modified the task so that the observer viewed only distance information. An example of this visual "phase-deaf" display is shown below (right) for a stimulus with 7 observations sampled from the signal distribution (the stimulus on the left is the identical sample with a spatial display including phase information). We used sample size $N = 20$ to roughly equate our task parameters to those used in the acoustic



experiment. An observed ROC curve for a well-practiced observer with this visual no-phase condition is shown as filled circles (dashed lines) in Figure 4. Note that performance in the visual modality is still superior, despite the removal of phase information. If the only shortcoming in the auditory task were an insensitivity to phase, we should expect these empirical curves to coincide. There are still unaccounted for processing losses in the auditory domain.

Some additional insights into phase-deaf processing in this task are discussed in the Appendix.

Beyond differences in phase-sensitivity, what other modality effects might be operating? As one of the Conference attendees suggested, the unexplained auditory deficit may be attributable to limitations in temporal integration in the auditory system. In the auditory task it is not possible to present all of the data simultaneously as was done in the visual display. Limitations of the auditory system in integrating information over a temporal range may correspond in the visual domain to the number of dots that the observer can integrate with brief presentations. Thus, sequential display of the visual data should further reduce performance; whether or not this manipulation can account for the remaining modality differences is an open question that merits empirical investigation.

Discussion

In summary, we devised a simple but rich surveillance task (that could easily be done by machine alone) in order to study human cognitive abilities and limitations. We found that, in the visual domain, humans generally (although imperfectly) follow the prescription of the ideal observer; and that human cognitive processes are quite versatile in adapting to variations in visual display representations. We identified an isomorphic acoustic task and observed that humans are markedly inferior in this modality--attributable in large part, we believe, to an inability to process relevant phase information.

As we have seen, auditory processing of the surveillance problem appears to have inherent deficiencies (phase-insensitivity) that may not be possible to overcome in the acoustic domain. By presenting the information visually, substantial gains in performance are achieved. Even a relatively degenerate visual representation lacking phase information produces better performance than observed in the auditory task (Figure 4). The more complete visual representations in Figure 1 provide a higher level of "aiding" that yields even better performance. Nevertheless, human performance is still not optimal.

Toward understanding the causes of this sub-optimality, we conducted some informal experiments that revealed human biases in computing the centroids: Some individuals tend to give undue "weight" to outlying points, while others tend to do the opposite. We have not pursued the nature of such biases in the various spatial displays in Figure 1. One expects that these stimuli might emphasize different characteristics of the data, and that they could be used to induce different biases in performance. But the minimal differences in performance with these stimuli (Figure 3) must temper these expectations.

Finally, we are particularly interested in the nature of the cognitive processes that are brought to bear on this generic problem. Our current plans are to concentrate more heavily on modeling the human operator (rather than his optimal counterpart) in an effort to reveal these processes and the sources of their sub-optimality. We also plan to address experimentally some of the factors underlying the efficacy and utilization of decision aids in this simple problem so that we may bring about more effective implementations of decision aids in the operational Navy.

Notes and References

1. Herchman, R. L., and Greitzer, F. L. Detecting a change in target location: A comparison of human and optimal performance. Proceedings of the 5th MIT/ONR Workshop on C3 Systems, 1982, LIDS-R-1267, 133-138.
2. Green, D. W., and Swets, J. A. Signal detection theory and psychophysics. New York: Wiley, 1966.
3. Watson, C. S., Rilling, M. E., and Bourbon, W. T. Receiver-operating characteristics determined by a mechanical analog to the rating scale. Journal of the Acoustical Society of America, 1964, 36, 283-288.
4. Jeffress, L. A. Stimulus-oriented approach to detection. Journal of the Acoustical Society of America, 1964, 36, 766-774.
5. Let E_s = signal energy and N_0 = noise power density. Then the equivalence is given by $E_s/N_0 = NR^2/2\sigma^2$.
6. This phase information is not to be confused with the unknown phase of the sinusoid signal.
7. We thank C. Rogers Saxon for helping with data collection and for providing important insights and criticisms during the preparation of this report.

Appendix:

"Optimal" Phase-Deaf Performance

If indeed human observers are unable to process phase information in the auditory task, then it is appropriate to compare their performance with an "optimal" processor who is similarly deprived. Such an analysis would also provide the theoretically optimum performance to which we can compare the visual no-phase data shown in Figure 4.

One obtains this theoretical processor in the same manner as in the original problem that we reported in [1]; i.e., derive the likelihood ratio of the density functions given states S_R and S_0 .

Given S_0 , the density function for each observation is circular normal when phase information is available. In the phase-deaf case, only the range r_i is given; in this case the observation r_i has the Rayleigh distribution,

$$f_0(r_i) = (r_i/\sigma^2) \exp(-r_i^2/2\sigma^2),$$

where the "0" in the subscript denotes state S_0 is given. The joint density of the N independent observations in the sample is, of course, the product of N such terms.

Given S_R , the range r_i for each observation has a Rice distribution,

$$f_R(r_i) = (r_i/\sigma^2) \exp[-(r_i^2 + R^2)/2\sigma^2] I_0(Rr_i/\sigma^2),$$

where $I_0(\cdot)$ is the modified Bessel function of the first kind and order zero. The joint density of the N independent observations is the product of N such terms.

The likelihood ratio is then

$$f_R/f_0 = \exp(-R^2/2\sigma^2) \prod_{i=1}^N I_0(R^2/\sigma^2),$$

and the decision requires a comparison of its value for the observed data with a constant $K = (c_R p_0 / c_0 p_R)$ given by the prior probabilities p_0 and p_R and the costs c_R (false alarm) and c_0 (miss); cf. [1]. Just as in the case of the phase-sensitive processor, this phase-deprived processor uses Bessel functions on distances to arrive at a decision. But in this case we were unable to simplify the expression. Therefore, we generated "optimal" no-phase performance by computer simulation of 10,000 trials to arrive at predicted ROC curves.

These optimal curves are shown in Figure 5 for sample sizes 1, 3, 5, 10, and 20; and parameters $R = 1.2$ and $\sigma^2 = 1.0$. The observed data from the visual no-phase task in Figure 4 are re-plotted in Figure 5 for comparison with optimal performance.

What can be inferred from the optimal models (both with and without phase) about the way that our observers process the information? Based on the theoretical analyses, it is clear that the no-phase condition is computationally more difficult than the phase-sensitive case. The phase-sensitive processor merely computes the sample statistic \bar{d} and compares it with the criterion d^* (which depends on the computation of a single Bessel function). A human observer in the

original task with phase information displayed (e.g., DOT5) does not actually need to compute Bessel functions; it is sufficient to compute \bar{d} and then base the decision on an estimate of d^* that is acquired through learning. On the other hand, the products of Bessel functions that arise in the no-phase case apparently do not reduce to such a simple comparison. It does not seem likely that humans could perform such complicated computations in their heads--in fact, we found that a simpler (but not optimal) rule based only on the mean of the sampled distances does nearly as well as the optimal (the ROC curves for this processor nearly coincide with those of Figure 5). Thus, as in the original task, the human decision maker can perform this task very adequately merely by computing means and learning through adaptation where to place the decision criterion.

It is also of interest to compare the optimal ROC curves in Figure 5 (without phase information) with those in Figure 2 (with phase information). First, note that for $N = 1$ both theoretical curves coincide, as we should expect since one observation does not provide any useful phase information. Second, we find that the no-phase curve for $N = 20$ virtually coincides with the phase-sensitive curve for $N = 7$. We interpret this to mean that the theoretical sensitivity for $N = 20$ (no phase) is the same as that for $N = 7$ (with phase). We might guess that human performance curves for these two conditions would also closely agree--and this is in fact the case, as a comparison of Figures 5 and 2 reveals. While this agreement is gratifying, it would be very useful to specify a functional relationship between the ROC curve sensitivities in the two conditions.




Figure 5. ROC curves for an optimal processor deprived of phase information; sample sizes $N = 1, 3, 5, 10$, and 20 . Each curve was generated by simulation of 10,000 trials. Plotted points are for a human observer with $N = 20$ (from Fig. 4).

THE EFFECTS OF A SPATIAL INFORMATION FORMAT ON DECISION MAKING PERFORMANCE IN
A C³ PROBABILISTIC INFORMATION INTEGRATION TASK.

Captain Brad Scott

Christopher D. Wickens

U.S. Military Academy

University of Illinois at Urbana-Champaign

SUMMARY

Eight subjects performed a probabilistic information integration task in which multiple cues, and their diagnostic value in choosing between the hypotheses. Primary interest focussed on the comparison of digital and analog-graphical displays, and upon sources of non-optimality in the information integration process. Generally the analog displays proved superior to the digital, and performance was better when information was presented at a more rapid rate. Subjects appeared to be non-optimal in their treatment of reliability vs. diagnosticity. However, departures from optimality related to serial position effects and to the use of the appropriate model of information integration were not observed.

INTRODUCTION

In modern command, control, and communication (C³) situations, the executive decision maker is often faced with the task of rapidly integrating a large number of sources of probabilistic information or cues, that bear on the likelihood that one or more competing hypotheses may be correct. The limits of human decision making in such situations have been amply documented (Wallsten, 1980; Slovic, Fischhoff, & Lichtenstein, 1977; Wickens, 1983). In the present report we consider the impact of three particular cognitive limitations in a simulated C³ scenario in which multiple sources of probabilistic information are to be integrated and one of two hypotheses are to be chosen.

(1) Sequential presentation. When probabilistic cues are presented sequentially over time there are two important ways in which humans have been found to depart from the optimal manner of integration. i) They have been found to give too much weight both to the first cues in the series (anchoring or primacy), as well as to the final cues (recency), in situations in which all cues should, optimally be provided equal weight. These are known as serial position effects. ii) In Bayesian decision making tasks, subjects often adopt "averaging models" of information integration in which information that weakly supports a hypothesis serves to reduce subjective confidence in that hypothesis, relative to the alternative. In fact, weak evidence should, optimally increase confidence, although by a lesser degree than strong evidence (Lopes, 1982).

(2) Reliability and diagnosticity. A given cue may be related probabilistically to a hypothesis in one of two ways. (i) Its diagnosticity (D), the probability of the hypothesis given the cue value, may be less than one. A completely non-diagnostic

cue (D = 0) may be a symptom that is equally likely under each of two hypotheses. (ii) Its reliability (R), the probability that the cue actually has that value, given the observed value, may be less than one. An unreliable cue value will have 0 correlation with the actual cue it purports to represent. When cues vary in both D and R the optimum decision maker should down weight the cue valence equivalently for both. Yet evidence suggests that humans do not always do so, but instead may treat probabilistic evidence as if it were totally reliable, when there is also variability in D (Johnson et al., 1973; Wickens, 1983). This bias will be referred to as the "as if" heuristic.

(3) Display format. The probabilistic values concerning R and D may be presented in digital or analog format. Experimental data suggest that precise digital readings may not be an effective way of integrating rapid information in order to obtain a "ballpark" analog estimate of some value (in this case confidence). The theory of stimulus/central processing/response (S-C-R) compatibility proposed by Wickens, Sandry, and Vidulich (1983) suggests that information requiring analog operations in working memory (such as the updating of a continuous scale of confidence) will be best served by graphical/pictorial display formats.

EXPERIMENT AND METHOD

In the present experiment our subjects participated in a tactical battlefield scenario. They were to imagine themselves as commanders of a defensive unit preparing to be attacked from one of two directions, North or South. In each decision problem, diagnostic information concerning the most likely direction of enemy attack was provided by a series of 6, 8, or 10 intelligence cues (varied between problems), presented at either a slow (5 sec/cue), or fast (3 sec/cue) rate. Each cue was identified by its source (i.e., air surveillance, ground topography), its diagnosticity (from 0 - 1.0), and its reliability (0 - 1.0). Problems differed in terms of the net amount of difference offered in favor of the most likely hypothesis, and this variable created three levels of support: weak, medium, and strong. Most importantly, for half of the problems, information was presented in the verbal format: a series of cues of the form shown in Figure 1. For the other half, the spatial format of Figure 2 was employed. One particular advantage of the spatial format, in which R and D defined the width and height of a rectangle respectively, lay in the fact that the total worth of valence of the cue is equal to the product of R x D. That is, the area of the rectangular cue. To the extent that area is a commodity that

is directly perceived, then the cognitively loading mental multiplication of $R \times D$ is avoided in the spatial display. The subject must simply integrate areas for one hypothesis or the other.




Figure 1: Information display with verbal code format.




Figure 2: Information display with spatial code format.

Eight male subjects performed in the experiment over a period of two - 1 1/2 hour sessions. Each subject received a total of 72 decision problems. This number was formed by the orthogonal combination of two display formats \times three levels of problem size \times three levels of net evidence \times two presentation speeds \times two levels of cue variability. This last variable manipulated the degree of correlation between R and D within a problem, but will not be discussed in detail here. Each problem consisted of an alternating sequence between cues in favor of the North and the South. On half of the trials the net evidence favored the North and on half the South was favored. Because the cues alternated, on half of the

problems the favored hypothesis was supported by the first cue and on half it was supported by the last. This difference was used to examine serial position effects.

RESULTS

Problem and display variables. Table 1 shows the mean accuracy of judgments as a function of each of the independent variables examined in isolation. There were no interactions between these variables and hence we present only the row and column means. Statistical analysis of these data revealed that the spatial display gave reliably more accurate judgments than the verbal ($T_7 = 3.5$, $p < .01$), that decision accuracy was lower at the slower presentation rate ($T_7 = 4.27$, $p < .01$), and was lower with more variable cue values ($T_7 = 2.17$, $p < .05$).

More precise information regarding the use of the probabilistic information was provided by the analysis of confidence ratings. These are shown in Figures 3 and 4. The actual ratings were transformed to a range of 0 - 20. On this scale, 10 indicates neutral confidence, 20 extreme confidence in the correct hypothesis. Values less than 10 occurred when subjects indicated confidence of varying degrees in the incorrect hypothesis. The confidence data are shown in Figures 3 and 4. Figure 3 depicts the strong effect of the three levels of absolute difference in evidence on the confidence rating ($F_{2, 14} = 25.1$, $p < .01$). These data are important in that they indicate that subjects performed the task appropriately, extracting more evidence from the data as more evidence was warranted. The importance of the linearity in these data will be discussed further below. Net evidence did not interact with any of the other variables.




Figure 3: Effect of net difference in evidence between the hypotheses (objective confidence) on subjective confidence.




Figure 4: Combined effects of problem size, display format, and presentation speed on subjective confidence.

Figure 4 presents the combined effects of three additional variables on the confidence ratings: problem size, display format, and presentation speed. The figure indicates that confidence was greater at the fast speed ($F_{1,7} = 5.88, p < .05$), particularly with the verbal format (F interaction = 6.34, $P < .05$). In addition, the figure reveals the interaction between presentation speed and problem size ($F_{2,7} = 5.97, p < .01$). Speed had little influence when the problems consisted of 6 or 8 cues, but when the problems were long, decision performance was lowered at the slow speed.

Optimality in information integration. Figure 3 reveals that subjects extracted more diagnostic information as progressively more information was offered. What makes these data particularly significant is the high degree of linearity between the objective and subjective confidence. Through analysis and model fitting described elsewhere (Scott & Wickens, 1983), we conclude that subjects are integrating the probabilistic information from the several cues in the optimal manner, rather than following the non-optimal averaging strategy.

A second analysis was performed to determine if there were marked serial position effects in the data, indicative of either anchoring or recency. This was done by comparing performance on trials when the correct hypothesis was supported by the first, and by the last cue in the problem. Neither accuracy nor confidence differed between these two sorts of trials, allowing us to conclude that either serial position effects did not occur, or if they did, then the primacy and recency effects perfectly balanced each other.

Finally, a second small experiment was conducted, using the same subjects, to determine the extent to which R and D may have been treated asymmetrically, with subjects showing biases to overweight low levels of R, applying the "as if" heuristic and treating those cues "as if" they were fully reliable. This determination was accomplished by constructing problems such that if the "as if" heuristic were employed, subjects would be induced to shift confidence toward the incorrect hypothesis, thereby producing errors

and lowering their net confidence rating. The results of this second experiment revealed that subjects were in fact non-optimal in this regard, overweighting cues with low levels of reliability.

DISCUSSION

The practical implications of the present data are two-fold. On the one hand, the differences in display format suggest an advantage to the spatial display for the kinds of judgments requested here. When information is presented at the relatively rapid rate, characteristic of both the slow and fast speed, decision makers cannot be expected to perform the necessary mental multiplication on the numerals of the verbal format in order to determine the aggregate evidence. The rectangular analog display provides this information in a more direct compatible format. The advantage of the spatial display is also supported by the interaction of display format and presentation speed on confidence. When the display is spatial, confidence is unaffected by speed. When it is verbal on the other hand, confidence is reduced at the slower speed, when the burdens on working memory are enhanced.

The second set of implications pertains less to engineering guidelines in system design, than to an appreciation of the kinds of cognitive limitations faced by the decision maker in the multi-element decision-making task. Recognizing that the effects observed here may be paradigm-specific, these may be briefly summarized as follows:

- 1) Information integration is limited by memory factors. When the problems are long (10 cues) and the display speed is slow, performance as measured by confidence, deteriorates.
- 2) Subjects appear to provide cues of low reliability with greater weight than is optimal.
- 3) On the brighter side, subjects appeared to use the appropriate model of information integration, eschewing an "averaging model" in favor of one that sums information for alternative hypotheses. Also there is little evidence in the present data for nonoptimal overweighting of either early (primacy) or late (recency) cues in the sequence.

While it is possible in the long run that engineering corrections may be implemented to circumvent those cognitive limitations of the second set that are encountered, it is at least important in the short run that their existence be acknowledged.

ACKNOWLEDGMENTS

This research was supported by Contract #N000-14-79-C-0658 from the Office of Naval Research Engineering Psychology Program. Gerald Malecki was the technical monitor.

REFERENCES

Johnson, E.M., Cavanagh, C., Spooner, R.L., & Samet, M.G. Utilization of reliability measurements in Bayesian inference: Models and human performance. IEEE Transactions on Reliability, 1973, 22, 176-183.

Lopes, L.L. Toward a procedural theory of judgment. Technical Report, Wisconsin Human Information Processing Program (WHIPP 17), Madison, WI, December 1982.

Slovic, P., Fischhoff, B., & Lichtenstein, S. Behavioral decision theory. Annual Review of Psychology, 1977, 28, 1-39.

Wallsten, T.S. Processes and models to describe choice and inference. In T.S. Wallsten (Ed.), Cognitive Processes in Choice and Decision Behavior. Hillsdale, NJ: Lawrence Erlbaum, 1980.

Wickens, C.D. Human Performance and Engineering Psychology. Columbus, OH: Charles Merrill, 1983.

Wickens, C.D., Sandry, D., & Vidulich, M. Compatibility and resource competition between modalities of input, output, and central processing: Testing a model of complex task performance. Human Factors, 1983, 25, 227-248.

Wickens, C.D. & Scott, B. A comparison of verbal and graphical information presentation in a complex information integration decision task. Engineering Psychology Technical Report EP-83-1/ONR-83-1, University of Illinois, Champaign, IL, June, 1983.

Table 1
Percent Accuracy

Trial Variability	Low	High		
	97.2%	92.3%		
Coding	Spatial	Verbal		
	97.6%	92.0%		
Time	Slow	Fast		
	93.4%	96.2%		
Weighted Difference	5-10%	15-20%	25-30%	
	94.2%	94.3%	94.8%	
Set Size	6(Total)	8(Total)	10(Total)	
	95.3%	95.3%	93.8%	

AD P002888

DEVELOPMENT OF A GENERALIZED HUMAN-MACHINE INTERFACE

Rita E. Knox

Lockheed Electronics Company, Inc.
1501 U.S. Highway 22
Plainfield, New Jersey 07061

Abstract

The development of a Generalized Human-Machine Interface is driven by consideration of human communication capabilities and limitations. The goal is to develop a system which provides machine capabilities similar to those required for communication among human beings. System features resulting from this approach and incorporated in the design include: application independence, attention monitoring, dynamic device assignment, human performance monitoring, and natural language processing. In addition, a special data management structure has been designed. System architecture and development progress are described.

Introduction

There is little disagreement that the human is a necessary and critical element in complex systems, and that poor performance on his part may impair the effectiveness of those systems. As acknowledgement of this, there has been a recent, renewed emphasis on 'Human Factors Engineering' and related issues (such as 'user-friendliness') in the design and development of new systems. However, the primary focus is often placed on the input/output devices to be used in the human-machine interface [1]; moreover, the findings of recent psychological research are often unknown to the applied community, where those basic results would have a significant impact on the design and development of new systems. A most distressing consequence of this situation is that many systems are designed and developed under circumstances where there is rarely sufficient time or money to evaluate design decisions, and those design decisions are only slightly influenced by empirical results that have been obtained under scientifically valid conditions.

I suggest that, while the input/output devices are essential components of the human-machine interface, they are not a sufficient description of that interface. In addition, although many properties of the human-machine interface that we consider important are present implicitly in the choice and application of devices, these more subtle features are often obscured by an obsession with the salient, physical characteristics of the hardware.

This paper will describe the design and development of a Generalized Human-Machine Interface (GHMI). This work has assumed, axiomatically, that if human communication with machines is to be successful, machines must 'know' something about human communication. That is, communication among humans depends on the sharing of common symbols, common assumptions, and common attributes of human cognitive abilities; since human communication is the 'model' of communication of greatest familiarity to humans, it is

the most intuitive model they can bring to communication with a machine; to the extent that the model is reflected (i.e., not violated) by the machine's activity, communication between human and machine will be simplified, efficient, and effective. The goal of the GHMI program is to develop a computer-based model of the structures of human communication.

System Overview

An overview of the GHMI is presented in Figure 1. The interface itself consists of two major components: an advanced console, and software which is the computer's 'communication intelligence.' The GHMI is truly an interface -- between the human being and a 'host system,' which is the 'rest of the system' of which a particular human and GHMI are only one subsystem. For example, the 'host system' might be a network of other humans and their workstations, a data acquisition and processing system, or a host computer in the usual sense.

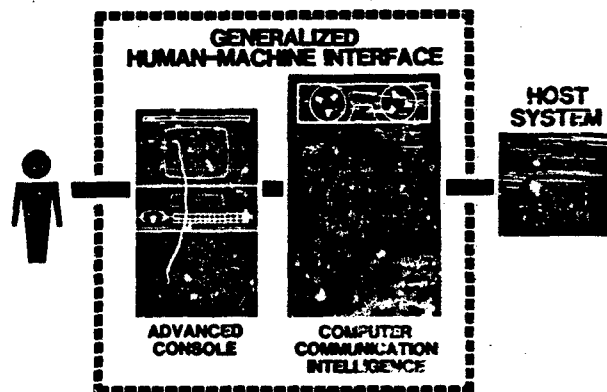


FIGURE 1

A specific 'advanced console' is only one example of many possible configurations. A particular console configuration will depend on the operations and tasks to be performed, the environment (e.g., land, air, or sea), and the platform (e.g., tank, living room, or helicopter). The 'advanced console' is a generic reference to advanced devices, such as voice recognizers and touch sensitive displays. These devices are advanced by virtue of their being easier for humans to use than the traditional buttons, knobs, switches, and keyboards.

The GHMI system architecture is 'generalized' in the sense that it neither depends on nor incorporates any

features of a specific application. This is accomplished through a complete separation of the structures of the system (i.e., the code) and the data required by a particular system. This design feature is not a simple consequence of adopting good programming techniques. Rather, it results quite naturally from our attempt to model human abilities, where many cognitive attributes and skills do not change as a function of the information being manipulated.

Of course, for this generalized HNI to be used in a particular system it must have access to the relevant data and constructs of that system. (By analogy, one must know about "meter", "movement", and "key", for a discussion of music as opposed to "rhyme scheme", "verse", "metaphor", and "alliteration" for a discussion of poetry.) Such information is made known to the GHNI through a process called "tailoring", as shown in Figure 2. The application-specific features introduced are the data used by the GHNI components (discussed below), and the console design, which consists of the choice and arrangement of input/output devices. As a by-product, this architecturally-driven approach to the development of an application-specific HNI allows economy in the cost and speed of new system development by having the "head-start" of a common, basic system.




FIGURE 2

GHNI Components

The computer components of the GHNI were designed to address specific human capabilities and features; the corresponding human and computer sites are shown in Figure 3. The fifth major component, the Data Base Interpreter and its associated Data Base, was not dictated primarily by consideration of human capabilities. Rather, it allows us to preserve the application independence of the system. These five major components are described below.

HUMAN CAPABILITY

- NATURAL LANGUAGE
- ATTENTION SHARING AND SWITCHING
- INTERSENSORY EQUIVALENCE
- HUMAN ERROR

COMPUTER INTELLIGENCE

- NATURAL LANGUAGE PROCESSING
- ATTENTION MONITORING AND SWITCHING
- DYNAMIC DEVICE ASSIGNMENT
- HUMAN PERFORMANCE MONITORING

FIGURE 3

Natural Language/Natural Language Processing

Natural language [2] is a primary, if not the most important method of human communication. As the concepts to be "discussed" by human and machine have become more complex, and computer use by the general population has increased, the necessity for natural language processing capabilities in computers has increased. The natural language capabilities of humans, however, are still far from being fully understood. Consequently, natural language processors for computers are marginally successful.

In developing natural language processing as a GHNI component we were concerned primarily with its inclusion as one feature of "communication intelligence;" the processor we've designed is not comparable to those developed in laboratories where natural language processing, *per se*, has been the primary focus. [3] Rather, the processor allows a certain amount of flexibility in the way a human can issue requests or enter data to the system.

The Natural Language Processor (NLP) is a special component of the more general Decode module which interprets all human inputs to the system. The translations made by the NLP are determined by the grammar which we designed. This grammar handles queries about the STATE [4] of an OBJECT's ATTRIBUTE; the comparison between the STATES of two OBJECTS, or between an OBJECT's STATE and some specified VALUE; the CONDITIONALS, IF, WHEN, WHENEVER; and the entry of an OBJECT's VALUE. The processing consists of "keyword" identification (i.e., lexical analysis), syntactical analysis, in which STATEMENT CLAUSES are synthesized by simple left-to-right processing rules, and semantic analysis, in which features like THING-ADJECTIVE and OBJECT-STATE agreement are assessed.

Attention Sharing and Switching/Attention Monitoring and Switching

Whether we consider communication between two people, or private thought, it is rarely, if ever, true that a "topic" is exhausted with no intervening digressions or diversions to other "topics." Intuitively, such multiplexing among topics appears to be a useful, necessary, and possibly unavoidable feature of human thought, resulting, at times, in creative insights drawn from consideration of seemingly unrelated ideas.

In performing tasks, either alone or with others, switching among topics is often necessary to understand a particular point, or is dictated by external events which must be attended to as they occur. The switch from one topic to another is rarely signalled by an explicit indication, and usually happens with no loss of general continuity. Moreover, a return to a deferred topic can often be accomplished without a major recapitulation of completed tasks or already discussed ideas; that is, the topic can be continued at the point where it was suspended.

A complete operation on some system will often consist of many more tasks than can reasonably be active at any particular moment. On the one hand, all tasks are not needed at all times; on the other hand, there are limits to a human's multiplexing abilities, thus making it, at best, superfluous and, at worst, confusing to have more tasks active than he can attend to. Thus, at any instant in an operation, only a subset of all tasks should be active; however, the number of tasks and their identities cannot be defined *a priori* since the particular subset is dictated by external events, and the chosen activities of a particular individual (e.g., an operator). [5]

The Attention Monitor holds the definitions of all TOPICs of an operation. These definitions consist of all TASKs of a TOPIC and the ordering constraints for their execution; the order operators specify completion, sequencing and repetition requirements for each TASK (and subsets of TASKs) within a TOPIC.

When the Attention Monitor receives an input (either from the human or the host system) it determines which TOPIC the input is associated with. If it is identified as a TASK of an ACTIVE TOPIC, checks for sequence and data range limit errors (if applicable) are made; if no errors are detected, the input, along with a TOPIC IDENTIFIER, is passed to the Data Base Interpreter. If an error is detected, the input is passed to the Human Performance Monitor.

If the input is not part of an ACTIVE TOPIC, it is a TASK of a NEW or DEFERRED (i.e., previously active, presently suspended) TOPIC. The Attention Monitor identifies the TOPIC and attempts to activate it by issuing a request to the Dynamic Device Assignment (DDA) module. That is, whether or not the TOPIC will be activated depends primarily on the availability of a device for that TOPIC. The necessary evaluations are made by the DDA module. If DDA determines that the TOPIC can be initiated, the sequencing structure (a data array) for that TOPIC is integrated into the Attention Monitor's dynamic TOPIC pool.

Intersensory Equivalence/Dynamic Device Assignment

The term "intersensory equivalence" is a general reference to the human ability to represent the same information in a number of different ways. Specifically, the representations may require the use of different sensory modalities. For example, one may speak a person's name, which requires the use of the speaker's vocal-linguistic apparatus and the listener's auditory system, or write it on a blackboard, which requires motor production and visual processing systems, and although the physical representations are quite different, the meaning derived in both cases is the same.

This "multi-modal" representation capability is used frequently, particularly when people are having problems communicating some idea (e.g., "Can you draw me a picture???"). Often we are forced to use a less than optimal representation because the modality of choice is not available under certain circumstances (e.g., when we're forced to use the telephone to give directions to some location). In general, while it may be the case that certain representations are better or worse for particular kinds of information, it is certainly not the case that there is a single modality (or a single representation within a modality) for the communication of a specific piece of information.

This human capability of intersensory equivalence implies, first of all, that alternative methods of presenting the same information should be a machine capability; also, it should be possible for the machine to know when the human has communicated the same information, even when different data representations have been used. When this concept is translated into a human-machine interface design it means that different devices (e.g., a display, voice synthesizer, etc.) must be capable of encoding or decoding the same piece of information.

In fact, such flexibility in a human-machine interface is essentially dictated by two concurrent developments: the transition from dedicated devices to general purpose devices, and the growing complexity of semi-automated operations. In practice, this means that a particular task may not always have access to

the device of choice if that device is being used by some other, possibly higher priority, task.

The Dynamic Device Assignment (DDA) component of the GHMI is responsible for generating the "optimal" configuration of TOPIC-to-(Virtual)Device assignments at each point in an operation when the set of active TOPICs changes. A virtual device may be identical to a physical device (e.g., a display unit), some portion of a physical device (e.g., a quadrant of a display), or to a cluster of portions of physical devices (e.g., two quadrants of a display and a voice synthesizer). For each TOPIC, we specify the set of virtual devices which can be used for display and control of the information of that TOPIC, along with the "suitability metric" for the use of each virtual device.

The DDA configuration processing is initiated by a request from the Attention Monitor to activate a previously inactive (i.e., new or deferred) TOPIC. The resulting configuration depends on (1) the number and identity of active and to-be-activated TOPICs, (2) the TOPIC priorities, (3) the available virtual devices, (4) the TOPIC-to-device suitability metrics, and (5) the "transition penalty" for moving an already active TOPIC from one virtual device to another.

This DDA processing may result in an active TOPIC being deferred, or in the request for activation of a new TOPIC being blocked. The Attention Monitor is informed of any changes in the set of active TOPICs and makes all necessary modifications to its TOPIC array. If any configuration changes are required as the result of DDA processing, this information is sent to the devices, where the new configuration is implemented.

Human Error/Human Performance Monitoring

Noise is an inherent property of all physical systems. In particular, noise in the human system, which sometimes results in human error, is an inescapable fact which may be reduced through training, education, or high motivation, but not eliminated. Specifically, it can never be designed out of a system without fully automating that system.

The danger of human error in existing systems results from the high credibility granted to any human input. Those inputs are propagated through systems, and, if they're inaccurate, the results may be disastrous unless some human detects and corrects them before their full impact is felt. Of course, if the correct inputs were known, a priori, the process could be automated and the requirement for the human eliminated. Most of the time, however, the human role in a system is to perform those tasks which we don't know how to automate. This presents a dilemma: How can a system detect an error if it doesn't know what the correct input should be?

A related issue concerns how well the human should (or must) perform in systems. System designers usually specify worst case limits on human task performance, such as maximum time allowed for a particular action or RMS error for some entry. These values are often arrived at by an arithmetic manipulation which takes the total system throughput requirement, subtracts the contribution of the non-human (i.e., designed) components, and specifies the residue as the tolerable limits on human performance. While some comparison of human requirements and human capabilities is made to ascertain whether the human will be able to perform as specified, the analyses are most often superficial. [5]

There are a number of issues to be considered if the GHMI is to reduce the frequency of human errors. First, the human component is not a system module

designed by human beings; thus, any predictions of human performance must come from models derived from observation of human behavior. A particular human's performance can be judged as good or bad only with respect to normative human performance; it can only be said to be desirable or undesirable with respect to any other, arbitrarily specified, design requirement.

Second, many human errors are not content errors; delay time in reaction to events, multiple corrections to a single input, and total time to construct an input are all associated with poor performance and could serve as data in the assessment of performance.

Third, any measure of human performance is a (possibly n-dimensional) random variable; therefore, decisions about performance must be based on statistical inference.

Fourth, the errors which can be detected by the machine must be caught at a point in the system which is near to the human who made them, before they are passed on to the system at large as a human input; moreover, an attempt to correct these errors should be made at the same local level as their detection.

Fifth, human performance monitoring having features of the preceding four points must be a continuous and (near) real-time HMI capability. Individual human performance will vary as a function of task, day, and time of day; performance evaluation on any crude time schedule will do little to reduce system failures attributable to human error.

A Human Performance Monitor (HPM) having the properties described above is summarized in Figure 4. At present, the HPM component implemented in the GHMI system has no assessment capability and only minimal measurement and feedback abilities. Design work for the complete module is in progress.

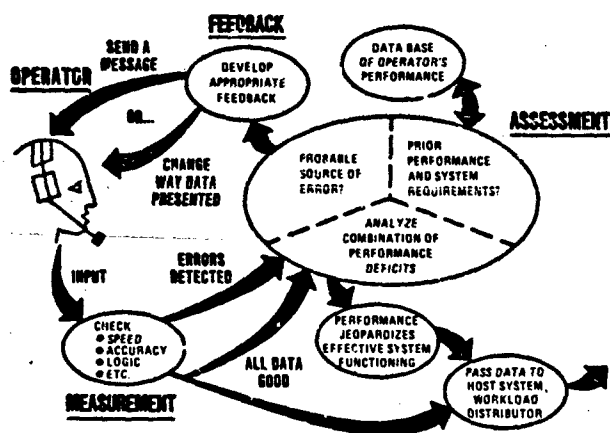


FIGURE 4

Currently, there are a number of errors that can be detected by the Natural Language Processor and the Attention Monitor. These are passed to the HPM module where one of two forms of feedback is given to the human. "Quick Mode", designed for experienced operators and urgent situations, immediately points out the operator's mistake to him. He is then directed to enter a corrected command, or in those cases where the system can guess at what the operator intended, he is asked to verify the system-corrected command. Upon positive verification, the command is executed.

"Teach Mode" is based on the philosophy that a human who discovers his own errors is less likely to repeat them than one who merely has his errors pointed out to him. Thus, when an inexperienced operator in a non-critical situation makes a mistake, he is taken through a series of successively more explicit hints about the identity of the error; after each message he is asked to correct the error if he can. If he cannot find his mistake after all hints have been supplied, the error is pointed out and he is asked to re-enter the command.

Data Base Interpreter

The Data Base Interpreter (DBI) provides a mechanism for storage and retrieval of information used by the HMI system. This information may come from the human or the host system, and may be rapidly changing, relatively static, or even a constant value, and may be continuously in demand or of only momentary relevance. The design of DBI was motivated by these attributes of the information and its storage and retrieval needs, and not by any explicitly human characteristics, as was the case for the other HMI modules.

Another consideration in the design of DBI was the philosophy of a general HMI architecture, demanding that any changes required by a new application be accomplished by changes of data, rather than by recoding of the HMI software itself. As a result of this approach, any specialized processing required by a specific application must be realized within the application-specific data base, and executed by the general-purpose Data Base Interpreter.

In the current implementation, this concern for separation of application-specific data from general-purpose code has been satisfied by implementing all application-specific processing in interpretive instructions which are stored with each data base object and are executed by DBI. This approach has the disadvantage of being inefficient for high volume computation or very large data sets; these demands could, however, be satisfied by other parts of the system outside the Human Machine Interface, or through execution of compiled subroutines by DBI, at the expense of pure software generality.

Two major advantages of employing interpretive instructions for application-specific processing are, first, the ease with which the system can be enhanced to satisfy new requirements and support new capabilities (e.g., new sensors or new sensor processing algorithms), and second, the efficiency with which operator tasks can be reallocated among HMI-based workstations, to satisfy the anticipated requirements of distributed processing systems, including distributed Command and Control. In the extreme, the role of an individual workstation could be completely redefined by the downloading of data.

The techniques used by DBI for storage and for retrieval are analogous to forward-chaining and backward-chaining, respectively. That is, when new information is provided to DBI, it is stored as the value of the appropriate data base object, and the values of other derived objects will also be automatically updated if they are listed as RECIPIENTS of the original object. This "forward-chaining" of updated values may continue to recipient objects outside of the database -- the human operator or the host system -- if one of them is to be informed of updated values.

When retrieval of information is requested, DBI may have the exact information already in storage, as a recipient of the automatic updating process described

above. Not all objects are handled in this way, however, since some objects are requested much less frequently than their component information is updated. For those objects, retrieval requires the derivation of the object's value from the stored values of lower-level DONOR objects. This "backward-chaining" retrieval of current values can, for some objects, result in requests for information being issued to the human or the host if one of them is the ultimate source of such information. Thus, DBI manages the flow of information between human and host by providing both automatic updating and retrieval of data upon request.

Figure 5 illustrates the architecture which ties these major GHMI "communication intelligence" components together, along with their supporting structures; the design is implemented in Pascal.

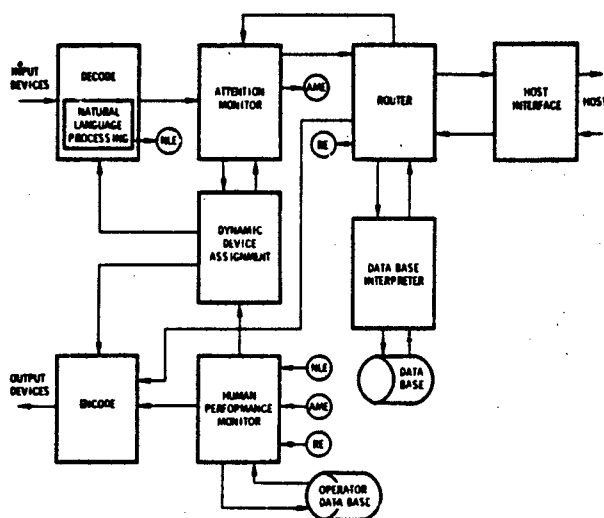


FIGURE 5

An Application

To test the HMI algorithms, a human must interact with them. For that to be possible, there must be communication devices and tasks for the human to perform. We have designed and fabricated a prototype console which houses two touch-sensitive high resolution displays (one color, one monochrome), a keyboard, and a voice synthesizer. This console serves as a testbed for the GHMI, and allows us to investigate the use of new input/output devices.

The operational tasks we've developed would be used in most command and control situations. At present, the land-based scenario includes map set-up, friendly, and enemy asset review and display, and network message review. Help functions, and alphanumeric readouts, which provide graphic symbol translation and status information, are also available. A photograph of one interactive screen display appears in Figure 6. At this time in the operation, the Friendly Assets and Enemy Assets Review TOPICS have been activated along with the situation display. The bar across the top of keys indicates an on-state: in the main menu at the bottom of the screen, it indicates the activation of topics; in the different TOPIC quadrants, it has the effect of plotting the designated units on the situation display. [7] The Dynamic Device Assignment module is responsible for arriving at this particular display/control configuration for the active topics.

The next TOPIC to be developed would include tools, such as unit position projection, for target correlation, fusion and aggregation.




FIGURE 6

The last piece of the application is the "host" in the system. In the present case, the host is a message network. A simulation advances 550 enemy tanks and armored personnel carriers in the area of interest; two line-of-sight sensors scan sectors of the battlefield and report number and position of sighted targets. Sensor messages, with the date-time group of sighting, are generated and sent to the HMI. These data form the basis of the analysis work performed by an operator.

The full system is implemented on two microprocessors, with the allocation of processing as shown in Figure 7. In the past we planned that the HMI software would run on a single processor; today we anticipate that multiple processors, possibly one dedicated to each module, will eventually be necessary to handle the high-speed processing requirements of the system. We have not begun to explore the implications of such a parallel architecture. The system has just recently become fully operational. Therefore, there are no data as yet on system or human performance.

GHMI's Role in Distributed Systems

As envisioned by many people, the architecture of distributed information processing and decision systems will require that resources be shared among individuals in the network, that collaborative analysis of situations be possible, and that the allocation of responsibilities to a particular individual be dynamically defined as a scenario develops. To this point, we have addressed a single node of a network, consisting of one human and the HMI system. What information could the HMI subsystem contribute to the full network?

The HMI has moment-to-moment knowledge of the tasks an individual is engaged in and of his current performance. In particular, it would recognize indications of overload or fatigue, or of highly efficient functioning. Such information on all individuals in a network could be sent to an independent, higher-level system module responsible for task allocation. With the addition of other resident knowledge (such as doctrine), this module would have the necessary information to dynamically define task assignments for each workstation so as to optimize

total system performance. (Notice the similarity, by analogy, to the GHMI Dynamic Device Assignment module; in the distributed system case, an individual is a "device" and his set of responsibilities is a "TOPIC.")

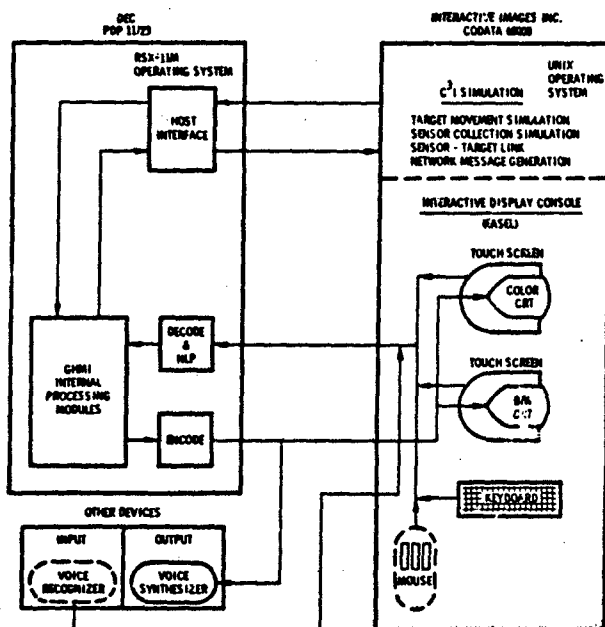


FIGURE 7

This network capability, summarized in Figure 8, is only one of many useful features of the GHMI. It was chosen for discussion since such a network function would not be realistic without the detailed knowledge of individual human performance and activity that can be accessed through the GHMI.




FIGURE 8

Conclusion

We have provided an introductory overview of the Generalized Human-Machine Interface designed and developed in our laboratory. We believe it to be a unique approach to the architecture of an interface between human and machine, one which should facilitate high-level communication. Our prototype, however, is just the beginning of what will be a much longer, but fascinating endeavor which will depend on the evolving understanding of human information processing and the revolution in computer technologies.

Acknowledgements

This work is supported as Independent Research and Development by Lockheed Electronics Company. In addition to the author, the following people are responsible for the design and development of the GHMI system: Richard Brandau, Merryll Herman, Ting Loo, and Joseph Schaul Jr.

References and Notes

- [1] See, for example, Huchingson, R.D. New Horizons for Human Factors in Design, McGraw-Hill, Inc., 1981; Schneiderman, B. Software Psychology: Human Factors in Computer and Information Systems, Winthrop Publishers, Inc., 1980.
- [2] By "natural language" we specifically mean written or spoken English. The grammar cannot handle German, for example.
- [3] See A. Barr and E.A. Feigenbaum (Eds), The Handbook of Artificial Intelligence, Vol. I, William Kaufman, Inc., 1981, pp. 281-316, for a review of more sophisticated natural language processors.
- [4] Upper case font indicates that a word is part of the formal terminology of the GHMI system. In the present case it indicates a part of speech of the Natural Language grammar.
- [5] Consider, for example, an operation consisting of 50 different tasks. A subset of active tasks is one element from the power set of the set of 50 tasks. Granted, many subsets may be eliminated as meaningless combinations or because the "reasonable number of active tasks" dictum is violated. However, the number of remaining configurations would make the description of an operation fairly complex.
- [6] Effectiveness of U.S. Forces can be Increased Through Improved Weapon System Design. G.A.O. Report No. PSAD-81-17, January 29, 1981, pp. 4-13, 27-34.
- [7] There are certainly other ways of designing interactive display layouts. We have emphasized the use of graphic symbology (the tactical symbols are based on FM 21-30) wherever possible. The on-off visual coding is really an arbitrary choice; in general, there are no guidelines for the design of such displays, and research in this area is needed.

TACTICAL DECISION AIDS WHICH QUANTIFY JUDGMENT

Rex V. Brown
and
Marvin S. Cohen

Decision Science Consortium, Inc.
7700 Leesburg Pike, Suite 421
Falls Church, Virginia 22043

Abstract

A new generation of decision aids, based on personalized decision analysis, is being developed to support the judgmental component of tactical command and control systems. They structure and quantify judgments of uncertainty and value, with a view to enhancing the soundness of decisions and reducing the decision maker's burden. Attack submarine examples are presented directed at: assessing target range; alerting to dangers and opportunities; and choosing a time of fire.

Problem: How to Support the Judgmental
Component of Tactical C²

Tactical command and control involves decisions on how to respond to changing awareness of dangers and opportunities, for example when to fire a weapon at a target. These decisions inevitably involve a blend of hard data (for example from sensors) and personal judgment (say, on the quality of sensor output). As tactical warfare and the threats it addresses become more complex, the judgmental component comes under increasing stress, particularly in the context of distributed decision making. Decision makers and their superiors find themselves in need of aids which will make the judgmental tasks easier to perform effectively.

Personalized Decision Analysis (PDA)
as a Promising Tool

The Office of Naval Research has funded a program of applied research to adapt the tools of personalized decision analysis (PDA) for this purpose. [1] PDA is a recently flourishing technique, grounded in statistical decision theory and engineering psychology. In principle, it can be used to enhance the quality of decisions, to reduce workload, and in some cases, to manipulate the command and control process. Its essence is to quantify the judgmental elements in a decision task and to display their action implications. In particular, uncertainty is measured by probability, value tradeoffs are measured by utility, and the preferred action is characterized as the one with the highest expected utility.

In 1979 a new effort was initiated by Engineering Psychology Programs of ONR through Decision Science Consortium, Inc., aimed at validating the practical value of this concept. [2] The intent was to foster in the fleet the actual implementation of at least one concrete variant. This was both to provide a stringent test of the applied merit of the concept and to uncover the most promising directions for further development.

Submarine ASW as a Testbed

Attack submarine command and control was selected as the testbed and within it the approach and attack ASW scenario. The initial thrust of the effort was to review the range of command and control activity in an attack submarine in order to identify situations of partic-

ular promise. In particular, we looked for command judgments where there was significant room for improvement, where any improvement would have a high impact on a mission, and where PDA appeared to have something to contribute. Three types of judgment were considered: assessment, alertment, and action selection. The situation identified for detailed attention was approaching an enemy submarine with intent to attack, which included each type of judgment. Specifically, we addressed the judgmental tasks of assessing target range (which became the primary initial focus of our project), alerting to critical situations, and choosing the time to fire a torpedo.

In each case, the object was to explore and develop a technology for aiding submarine commanders and their staffs in making these judgments; specifically to help them to assimilate the data and the expertise available to them and to combine them effectively, with contributions from their own judgment.

Assessing Target Range

The judgmental task chosen for our primary effort was the assessment of target range (including a margin of uncertainty) in the light of multiple conflicting estimates. A logical algorithm for this task has been developed with support from the Mathematical Sciences Division of ONR. [3,4] It takes individual solutions (estimates) and assessments of their accuracy (and interdependencies) and generates a pooled estimate with a band of uncertainty changing over time (Figure 1).




Figure 1. Illustrative Range Pooling Display

The source of the accuracy estimates may be personal judgment or previous analysis (say, by fitting parameters to exercise data). The algorithm appears well

enough established now to permit primary attention to be devoted to user interface issues. An early version of the algorithm showed substantial improvements over individual solutions and over the "system solution" when tested on fleet exercise (Rangex) data.[5] Work on engineering psychology aspects of eliciting judgmental inputs and presenting outputs is being supported by the Psychological Sciences Division of ONR.[6,7]

An interactive graphic computer program has been developed based on the responses of subjects representative of ultimate users, in situations which approximate realistic engagements. Figure 1 shows a typical display in the context of which the user can change the bias adjustment (shown as a wiggly bar) or the 80% limits, and see displayed the resulting change in the pooled assessment (shown at right). Figure 2 shows another form of desired output, in which pooled assessments are plotted against time. Implementation issues (including interface options) have been proposed for investigation in an experimental testbed in collaboration with a Navy laboratory (NUSC). Subsidiary PDA aids for assessing error in specific range estimates as a function of several sources of error, have also been explored.




Figure 2. Time-Range Display

Planning Attack and Determining Time-of-Fire

In addition to target range pooling, other aids supporting the time-of-fire decision have been addressed--for assessing the probability of hitting the target, alerting to critical situations, as well as actually recommending a time-of-fire. These aids are, however, at an earlier stage of development: the logical algorithms have been identified; the form of user interface has been anticipated; and further development is being supported by IBM Federal Systems Division.

Figure 3 shows the display format of an alerting aid. It monitors range-related probabilities such as being within a certain target's counterdetection range,




Figure 3. Alerting Aid: Output

of the target's being within own ship's weapon range, and range accuracy being within, say, 2000 yards. The aid emits a signal when critical thresholds of probability (which may be preset or set by the user) are reached (shown as broken lines in display). Assessments from which the probabilities are derived are shown on the right to permit validation by the user. (Note that they all involve range accuracy, which would be output from the range pooling aid, discussed above.)

The time-of-fire decision (how long should I wait, attempting to remain undetected, before launching an attack?) involves both uncertainty and value. There are the competing values of killing the target and preserving own ship, and uncertainties about the consequences of any given course of action. Timing is of the essence: A premature attack may both alert the enemy and miss; yet delay increases the risk that the SSN will be detected before attack or that the opportunity to attack will slip away. For the time-of-fire decision aid, all these considerations are put together into a "decision tree," described in detail in [6].

Figure 4 shows the basic features of the time-of-fire decision aid. It predicts the relative value of waiting to shoot at various points in time from the present, given a choice of own ship maneuvers. The point at which relative value is maximized is the recommended time, all things considered, at which shooting has the highest expected payoff. For purposes of comparison, the straight horizontal line, marked "Retire now," is the value of simply leaving the scene. (As the approach scenarios proceed, of course, the display would be updated to reflect new information.) Such a display would have two functions: to help decide now whether or not to fire, and, if the decision is to wait, to help make preparations for a future firing.

TIME: 1405




Figure 4. Attack Planning Aid: Output Display

A critical issue, currently being explored, is how best to provide these inputs in a way that does not put an unacceptable burden on submarine command staff in the heat of battle. The intention is to have the displays largely driven by more primary input data generated routinely on-board or from the results of previous analysis.

A Special Role for PDA in Distributed C²

PDA may have a particularly promising role to play in distributed tactical C² systems. Decision-analytic aids can influence the decisions of subordinate commanders in the system so as to balance autonomy from and control by higher authorities.

For example, an aid to indicate when a submarine weapon should be fired, such as we have discussed, would incorporate assessments of uncertainty (for example,

probability of kill and probable target range) and value judgments (for example, the relative importance of killing and avoiding being killed). The inputs to the aid may be specified ahead of time by higher authority (for example, value tradeoffs and the assessment of threat characteristics and sensor performance); or they can be provided on the spot by the aid user (say, an assessment of target intent). The aid can be used centrally by a Battle Group commander or by a subordinate commander. The output of the aid can be binding on the subordinate or merely suggested guidance.

The degree of centralized control can be determined by how the use of this aid is specified along such dimensions, for example by discretion in specifying inputs or in acting on output. That control can be changed by contingency rules on the aid's use, for example to control the subordinate commander tightly under routine circumstances, but to give increased autonomy if communication with Battle Group command should be broken. Even when the quantification and use of the aid is fully delegated to a subordinate commander, it may provide a reviewable record of his decision process, which may influence him to act in closer accordance with higher-level policy (as reflected in default inputs to the aid).

Figure 5 illustrates how the aid might be used to affect the degree of centralization of control. In the decentralized mode (hatched and solid lines) the submarine commander is informed by the aid but makes his own decision. In the centralized mode (hatched and dotted lines) he is required to follow the indication of the aid. (Note that many other variations are possible by changing the hatched lines, for example by giving the local commander a say in the value inputs, or the central commander a say in the uncertainty inputs.)

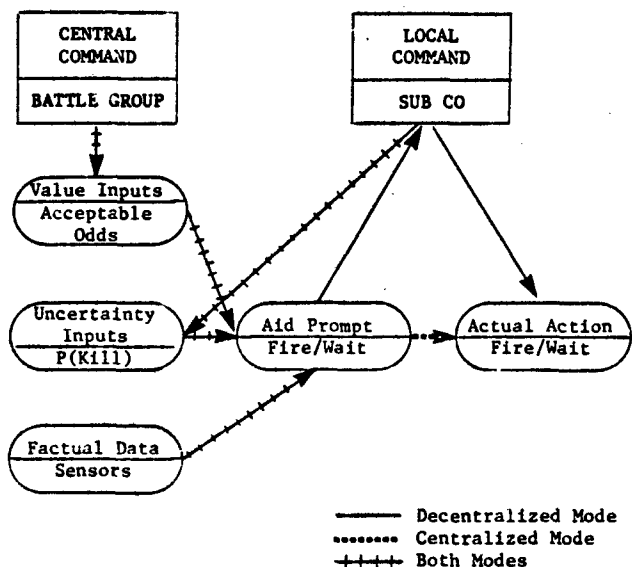


Figure 5. How Aid Might Affect Centralization

Conclusion

We have argued that there may be an increasing role for tactical decision aids which quantify judgment. The danger, of course, is that they may demand more judgment in their input than they save in the judgmental task they are designed to aid. This is why we attach so much importance to the human engineering aspects of aid development. Solving the people problems is more important than solving the logic problems.

Acknowledgements

The work described was supported largely by the Office of Naval Research, Engineering Psychology Programs, under contract number N00014-80-C-0046 (technical monitors M.A. Tolcott and G.S. Malecki). Some algorithmic development has been supported by ONR, Mathematical Sciences Division (technical monitor J.R. Simpson). Current work on the time-of-fire aid is being supported by IBM Federal Systems Division.

References

- [1] Brown, R.V., Peterson, C.R., Shawcross, W.H., and Ulvila, J.W. Decision analysis as an element in an operational decision aiding system (Phase II) (Technical Report 75-13). McLean, VA: Decisions and Designs, Inc., November 1975. (NTIS No. AD A018109)
- [2] Cohen, M.S., and Brown, R.V. Decision support for attack submarine commanders (Technical Report 80-11). Falls Church, VA: Decision Science Consortium, Inc., October 1980.
- [3] Brown, R.V., and Lindley, D.V. Improving judgment by reconciling incoherence. Theory and Decision, 1982, 14, 113-132.
- [4] Lindley, D.V. The reconciliation of probability distributions. Operations Research Society of America. September-October 1983.
- [5] Conen, M.S. Quantitative validity of objective and subjective inputs for a target range pooling decision aid. Proceedings of the 49th Military Operations Research Symposium (U) (MORS-49). Alexandria, VA: Military Operations Research Society, Inc., June 1982. (S)
- [6] Cohen, M.S. Decision support for attack submarine commanders: Target range pooling and attack planning (U) (Technical Report 82-1). Falls Church, VA: Decision Science Consortium, Inc., April 1982. CONFIDENTIAL
- [7] Bromage, R.C., Brown, R.V., Chinnis, J.O., Jr., Cohen, M.S., and Ulvila, J.W. Decision aids for submarine command and control. Phase III: Concept implementation (U) (Draft Technical Report 83-2). Falls Church, VA: Decision Science Consortium, Inc., March 1983. CONFIDENTIAL

A KNOWLEDGE BASED INTERACTIVE PROCEDURE FOR PLANNING AND
DECISION SUPPORT UNDER UNCERTAINTY AND PARAMETER IMPRECISION

Andrew P. Sage
Chelsea C. White, III

Department of Systems Engineering
University of Virginia
Charlottesville, Virginia

This article Abstract

We summarize key features of an interactive planning and decision support process for multiple criteria alternative selection situations. Probabilities, utility scores for the lowest level attributes, and attribute tradeoff weights, i.e., the parameters, can be imprecisely described by set inclusion. Within a specified structural model of the decision situation, the process allows the decisionmaker to iteratively select the mix of parameter value precision and alternative ranking specificity. By selecting this mix, the decisionmaker is able to direct the alternative selection process in an interactive manner, using alternative selection strategies based on behaviorally meaningful dominance search strategies. Emphasis is placed on the motivation of the research and the behavioral relevance of the support process.

1. Introduction

The process of choosing among multiattributed alternatives often involves an initial search for a dominance structure and ultimate identification of a set of nondominated alternatives. By definition, a nondominant alternative is one which is not worse than any other alternative on any attribute and which is better than each other alternative on at least one attribute. In most decision situations, however, there is no single alternative that dominates all other alternatives, at least initially. In such decision situations, the decisionmaker typically "adjusts" the structure of the decision situation, and parameter values within this structure, so as to identify a dominance structure which contains a single nondominant alternative. This search may involve rational activities, such as aggregation of attributes and compensatory tradeoffs through determination of judgmental weights. Alternately, it may involve various rules which may be quite flawed. Examples of such rules are i) lexicographic ordering, in which the best alternative on the most important attribute is selected, and ii) sequential pairwise comparison of alternatives using a preference relation that is a function of the two alternatives being compared. In this latter case, intransitive preferences may easily result due to the fact that the contextual relation used to determine preferences changes from binary comparison to binary comparison.

A variety of holistic, heuristic, and wholistic judgmental activities will typically be involved in the search for a dominance structure among the alternatives. These take on various forms and mixtures of formal knowledge based, rule based or skill based

activities as deemed appropriate for the task at hand [1,2]. Especially when there are a large number of alternative courses of action under consideration, the decision process will typically involve mixed scanning, where some noncompensatory rule is first used to eliminate grossly inappropriate alternatives. This is then followed by one or more compensatory information evaluation operations that results in a dominance structure which enables final judgment and alternative selection.

The research discussed here is based upon the hypothesis that people are able to evaluate alternative plans and decisions efficiently and effectively, and with low stress, when there is a clear dominance pattern among alternatives that enables establishment of a sufficiently discriminatory priority structure. Our goal is to provide a knowledge based decision support process that enhances the quality of the dominance structure used for judgment and choice.

The next section will present a summary discussion of the features and structural constructs of our decision support system. The following section presents a more detailed discussion of these structural constructs and introduces some of the modes in which the support process can be used. Then we discuss some behavioral issues that relate to the conceptual design of ARIADNE. The list of references contains citations to a number of works which discuss the algorithmic content of the decision support system.

2. Features of the Decision Support System

We now investigate concepts for the design and evaluation of an interactive knowledge based planning and decision support system which combines, or allows combination of, several evaluation rules and contingency structures often used as a basis for evaluation, prioritization, judgment, and choice. We have developed a knowledge based system to interactively aid planning and decision support processes through encouragement of search for a dominance structure that is behaviorally realistic and rational, from both a substantive and procedural viewpoint. The support system is called ARIADNE for Alternative Ranking Interactive Aid based on Dominance structural information Elicitation. The support system enables use of various integrated forms of wholistic, heuristic, and holistic reasoning in an aided search for dominance information among identified alternatives. We believe it to be flexible enough to closely match diverse decision situations and environments in order to support varying cognitive skills and decision styles, thereby enabling planners and decision makers to adapt its use to their own cognitive skills, decision styles, and knowledge.

Our efforts have concerned choice making situations under certainty and under risk, primarily for the single decision node case. This formulation allows

This research was supported by the Office of Naval Research under Contract N00014-80-C-0542.

consideration of a variety of imprecisely known parameters such as: attribute tradeoff weights, outcome state values on lowest level attributes, event outcome probabilities, and various combinations of these. Parameter needs are determined from the structure of the decision situation, as elicited from the decisionmaker during the formulation and analysis steps of the decision support process. We consider these formulation and analysis steps to be outside the scope of our present software developments but recognize the essential need for them in a complete decision support process.

The decision situation structural model may represent decisions under risk or under certainty. The attribute tree representing the features of decision outcome states may be structured and/or parameterized in a top-down or bottom-up fashion through use of ARIADNE. A single level structure or a multiple level hierarchical structure of attributes may be used with the choice of these being at the discretion of the decisionmaker. Multiple decision node situations may be approached through a goal directed decision structuring approach in which the growth of the structure of alternative decisions and event outcomes is guided by sensitivity-like computations obtained through use of the ARIADNE algorithms [3-5].

Parameters are elicited from the decisionmaker in the form of equalities and inequality bounds. A variety of mathematical programming approaches and graph theory, have been used to generate interactive displays of preference digraphs. These mathematical programming approaches are used to determine dominance structures for alternative prioritization that are based on parameter information elicited from the decisionmaker. At present, only a linear programming approach will yield necessary and sufficient conditions for determination of a priority structure and computational times that are consistent with interactive decision aiding. This requires that we elicit structural parameter information in a slightly restricted form which we denote the "behaviorally consistent information set" (BCIS). Often this BCIS will be in such a form that solution of the generally nonlinear programming problems associated with determination of dominance structures can be replaced by the solution of simple, computationally amenable linear programs with bounded variables. The major simplification associated with eliciting parameter imprecision in a prespecified structural format, however, is in the natural language dialogue needed to establish a model of the decision situation.

The purpose of the graph theory algorithms is to enable construction of a domination digraph, or dominance structural model. This digraph is a pictorial representation of the ordinal preferences as determined from a dominance reachability matrix. This matrix is determined by the linear programming algorithms from the decision situation structural model and parameters elicited from the decisionmaker. These domination digraphs encourage either selection of a preferred alternative, or further iteration using the aggregated preference information for feedback learning.

An inverse aiding feature is currently being incorporated into the decision support system [6,7]. This feature allows the decisionmaker to make wholistic, skill based prioritizations among alternatives. These prioritizations may be across some, or all, identified alternatives, at the top level of the hierarchy of attributes or at some intermediate level. If we elicit numerical bounds on the attribute scores for those attributes which are subordinate to and included within the attribute at which alternatives are prior-

itized, then bounds on attribute weights, consistent with the wholistic prioritization, may be determined by using a linear programming approach. Alternately, if weights are specified, then it is possible to determine bounds on alternative scores on those attributes subordinate to the attribute at which prioritization was made through use of linear programming algorithms.

As alternatives are identified and prioritized, updates on these bounds are made available. The results obtained from using the inverse aiding feature are, in many ways, comparable to those obtained from the regression analysis based Social Judgement Theory [8]. This approach provides weight identification only, with a "confidence" measurement concerning the validity of weights, cardinal preferences are assumed. Results in the form of bounds on, or ranges of, weights are available with a very few alternative prioritizations in the inverse aiding approach. The prioritizations needed may involve a mixture of cardinal and ordinal preferences. For a large number of prioritizations, the inverse aiding approach may become cumbersome computationally compared to the regression based approach, where additional information may be easily processed in a sequential fashion.

Combination of inverse and direct aiding to enhance decisionmaker specification of imprecise values, weights, and probabilities enhances the usefulness of ARIADNE since it allows for judgments and their explanation, using a combination of formal knowledge based and skill based modes. This enhanced usefulness will also occur through encouragement to the decisionmaker to become more aware of relevant alternative courses of action and to identify new alternatives on the basis of feedback learning of the impacts of alternatives upon issues and objectives in a behaviorally relevant way that, hopefully, encourages "double-loop learning" [9].

3. Structure of ARIADNE

A complete set of activities envisioned in using the single stage, or single decision node, version of ARIADNE involves the following set of activities.

Formulation of the Decision Situation

1. Define the problem or issue that requires planning and decisionmaking by identification of its elements in terms of
 - (a) Needs, and
 - (b) Constraints or bounds on the issue.

2. Identify a value system with which to evaluate alternative courses of action, and identify objectives or attributes of the outcomes of possible decisions or alternative courses of action.

3. Identify possible alternative courses of action, or option generation.

Analysis of the Decision Situation

1. Determine outcome scenarios.
2. Identify decision structural model elements, that is those elements or factors from the conceptual formulation framework which appear pertinent for incorporation into a decision situation structural model.

1. Structure decision model elements:

- (a) Structure decision tree,
 - (b) Structure information acquisition and processing tree--which may be part of the basic decision tree, and
 - (c) Structure attribute tree or objectives hierarchy.
4. Determine independence conditions among elements of the attribute tree and decision alternatives.
 5. Identify potential for the use of deficient information processing heuristics and provide appropriate debiasing procedures.
 6. Determine impacts of, or outcomes that may result from, alternative courses of action.
 7. Encode uncertainty elements in the form of event outcome probabilities, or bounds on these, to the extent possible.
 8. Identify risk aversion coefficients, if needed, to the extent possible.
 9. Identify preference or value functions, or bounds on these functions, to the extent possible.
 10. Identify attribute weights, or bounds on these functions, to the extent possible.
 11. Identify wholistic preferences among alternatives to the extent that this is possible.
 12. Identify possible disjunctive and conjunctive aspects, e. thresholds for attributes, of identified alternative courses of action.

Evaluation and Interpretation of the Outcome of Alternative Courses of Action

1. Identify a decision aiding protocol, or plan, for evaluation and interpretation of the decision situation.
2. Identify potential for use of deficient judgment heuristics.
3. Use conjunctive and/or disjunctive scanning to eliminate very deficient alternatives and retain alternatives meeting minimum acceptability criteria across attributes.
4. Determine the maximum amount of domination information possible:
 - (a) Display domination digraph.
 - (b) Identify alternative courses of action which could not be among the N most preferred alternatives. Normally these are deleted from further consideration.
 - (c) If the decisionmaker can select an alternative for implementation by wholistic judgment, or prioritize the remaining alternative set through heuristic elimination, then go to step 6 of Evaluation and Interpretation.
 - (d) If a choice cannot be made, then assess further information about values of imprecisely known parameters by iterating through steps 6-11 of Analysis (B). then return to

step 1 of the Evaluation and Interpretation (C). There exists many possibilities for obtaining greater alternative evaluation specificity such as:

- (i) setting higher aspiration levels or aspects,
- (ii) moving up the attribute tree by determination of a subset of attribute trade-off weights,
- (iii) "tightening" bounds on attribute trade-off weights,
- (iv) "tightening" bounds on event outcome probabilities, possibly through information processing updates,
- (v) "tightening bounds" on value or preference functions.

5. If the decisionmaker has provided (partial) wholistic preferences as part of the analysis effort, use these with the inverse aiding feature of the aid to determine bounds on attribute weights implied by these preferences such as to provide learning feedback to decisionmaker.
6. Conduct sensitivity analysis. Provide the decisionmaker with an indication of how sensitive the optimal action alternative, or prioritization of alternatives, is with respect to changes in values and information about impacts.
7. Evaluate validity and veracity of the approach. Encourage judgment concerning whether the formulation, analysis, and interpretation are sound. If not, encourage appropriate modification to structure and parameters associated with the decision situation, including identification of additional attributes and alternative courses of action. Then, iterate back to an appropriate step and continue.

In our work to date, we assume that the details of issue formulation and analysis are accomplished external to the interactive aid itself. There are a variety of procedures for accomplishing these tasks. [10] Our research assumes that there exists an issue formulation structure and that the impacts of alternatives are known. These are provided through various elicitation activities. We do not envision that the software we develop for interactive interpretation, including evaluation and prioritization, will generally be suitable for use independent of a trained decision analyst. Whether software can be evolved to result in an appropriate "stand alone" aid is very dependent upon the environment and other factors that constitute the contingency task structure for a specific situation. In situations which are repetitive and environments which are stable, such as in health care or equipment fault diagnosis situations for example, it seems entirely possible to design useful "stand alone" aids. In most strategic, and in many tactical situations there will not be a stable underlying structure that will easily allow this. The activities involved in issue framing and the identification of a dominance structure appropriate for decisionmaking are often very situation dependent.

There are a number of considerations that influence planning and decision support processes. The person using a decision support system should be aware of these considerations if best use of the aiding process is to be obtained. Generally these considera-

tions involve the operational environment and the familiarity of the decisionmaker with the environment and task at hand. It is the interaction of these factors that influence:

- (1) behavioral characteristics of the decisionmaker,
- (2) interaction between decisionmaker and analyst,
- (3) choice of computer-based support for decisionmaker analyst interaction

Among the behavioral characteristics of the decisionmaker that influence aiding consideration strongly are the facts that the decisionmaker:

- (1) is often impatient with time consuming and stressful assessment procedures;
- (2) wants to see some preliminary results promptly if these are needed or wanted;
- (3) may lack interest in interacting directly with complex quantitative procedures for decision aiding that do not seem tailored to the specific contingency task structure of the issue at hand; and, as a consequence,
- (4) requires a decision aiding approach that adapts to the decisionmaking style appropriate for the decisionmaker in the given contingency task structure.

There are a number of considerations that influence the most desirable interaction between the decisionmaker and the analyst. The interaction must be such that these result:

- (a) a list of objectives and an objectives hierarchy;
- (b) a list of alternatives; and
- (c) a list of outcomes for each alternative.

The extent of the need for the use of these identified lists will vary greatly with the "expertise" of the decisionmaker. A major task of the analyst in the formulation and analysis portion of the aiding effort is to assist the decisionmaker in obtaining these "lists" in a behaviorally relevant and realistic manner.

The analyst must also ensure, to the extent possible, that:

- (1) the above lists are reasonably complete;
- (2) the lowest level objectives are additively independent;
- (3) the alternatives are mutually exclusive, and
- (4) the outcomes that follow from each alternative are mutually exclusive and exhaustive.

The nature of the interactive process is such that iterative changes can be made in terms of addition or deletion of alternatives and attributes. Nevertheless, there are significant advantages in attempting to be reasonably complete at the start of the interpretation portion of the process.

The decisionmaker must provide the analyst, following behaviorally realistic elicitation procedures, information regarding:

- (1) alternative scores on lowest level attributes,
- (2) tradeoff weights,
- (3) probabilities, and
- (4) relative risk aversion coefficients

or appropriate ratios or bounds on these quantities which represent the precision that the decisionmaker believes appropriate or is capable of providing for the given decision situation.

There are many computer based support considerations which evolve from decisionmaker-analyst interaction considerations. A goal of all decision support system design efforts is to obtain "friendly software", software that is friendly both to the decisionmaker and the analyst. In particular, the analyst must be able to interpret the decisionmaker's structural and parameter information for input to the computer. To do this may require:

- (1) redefining the outcome space, such as redefinition of attributes to ensure satisfaction of independence considerations and
- (2) describing parameter information in terms of inequalities (or more generally set membership).

The analyst must be able to interpret computer output in a fashion that facilitates the decisionmaker's understanding and decisionmaking abilities. The analyst must be able to assist the decisionmaker in responding to the following question which is central in our interactive knowledge based support system:

Has sufficient preference and structural information been elicited from and provided to the decisionmaker for alternative selection, or is more information required for identification of a dominance structure that is relevant and appropriate for quality decision support?

If the decisionmaker feels that an alternative can be selected for action implementation at any stage in the interactive aiding effort, the analyst must be able to encourage decisionmaker judgment concerning whether or not the issue formulation, analysis, and interpretation are sound. If the issue formulation, analysis, and/or interpretation are not perceived as sound by the decisionmaker, the analyst must be able to encourage appropriate structural and parameter value modification, typically by means of sensitivity analysis, in order to insure effective, explicable, and valid planning and decision support. If the decisionmaker cannot choose an alternative from among those considered, the analyst must be capable of eliciting further structural and/or parameter information to enhance appropriate selection of alternative courses of action.

One very important feature of a knowledge based system for planning and decision support is encouragement to the decisionmaker for generating new options, outcomes, and attributes at essentially any point in the aiding effort; and ability to properly evaluate these new options. The analyst must be able to cope with this additional information under the assumption that:

- (1) the new information is consistent with previously obtained information; or
- (2) the new information is not consistent with previously obtained information due to
 - (a) structural inconsistencies, or
 - (b) parameter inconsistencies.

Thus, the capacity to resolve potential inconsistencies through interaction with the decisionmaker must exist within the planning and decision support process. The indirect, or inverse decision aiding, feature should be of particular value to this end. In a "policy capture" like fashion, this indirect feature will allow identification of bounds on attribute weights in terms of wholistic preferences among some, or all, alternatives. In the direct aiding feature, values, weights, and probabilities are identified and prioritization of alternatives result from this. Combined use of the direct aiding feature with indirect aiding should result in much learning feedback concerning relations among the various modes of judgment.

ARIADNE, as we have noted, does not contain software to assist in the formulation and analysis portion of the planning and decision support effort. It is in these two steps that alternative choices, attributes and decision impacts or outcomes are elicited or identified. Our effort is much more concerned with the interpretation part of a decisionmaking effort; that is to say how information is processed concerning formulation and analysis based quantities such as probabilities, values, weights, ratios, and bounds upon these. We are concerned also with the way in which this information is aggregated, by any of a variety of formal knowledge, rule based, or skill based modes of cognition that result in judgment and choice. We recognize the difficulties in separating the tasks of formulation and analysis from those of interpretation. There are difficulties at the systems management level since the way in which people cognize a problem, as part of the contingency task structure of a particular situation, determines the way in which they will go about resolving it. Thus the performance objectives, information processing style, and decision style that are most appropriate and that are likely to be used for a given task, are very much dependent upon the task itself. When a particular concrete operational or skill based strategy has yielded previous satisfactory results, many people will tend to use that strategy unquestioningly and uncritically in new situations perceived to be similar. This can result in very unsatisfactory judgments and choices in decision situations that have changed and that are not recognized as different from familiar past situations. This may result in premature cessation of search and evaluation of alternatives prior to identification of quality strategies, even for familiar situations. The efforts can be devastating in unfamiliar environments that are not so recognized [11].

The strategies which a decisionmaker will desire to use for interactive interpretation will be strongly dependent upon the way in which the task requirements are initially cognized. This will influence the objectives, attributes, and alternatives generated in the formulation step and the value scores or impacts associated with them in the analysis step. The input information to the interpretation step is just this information. Adequacy of the interactive interpretation step will clearly be dependent upon the "quality" of the information input to it.

Many recent studies [12] have indicated that people often construct selectively perceived simple

deterministic representations of decision situations that make information processing easy and which do not reflect the complexities and uncertainties that are associated with the actual situation. A goal of a decision support system is to encourage wide scope perceptions and associated information processing. The process used to assess probabilities, utilities, and weights will doubtlessly affect the quantities that are elicited. It is possible, for example, that a poor elicitation procedure may, unknowingly or knowingly, create rather than measure values [13]. An advantage to formal support for planning and decisionmaking processes is that it is possible to conduct a search for inconsistent judgment and perhaps even detect flawed information processing heuristics if process tracing is used. When inconsistencies are discovered, it then becomes possible, at least in principle, to examine the judgment process to determine which judgments imply flawed information processing, and/or incoherent or labile values, and/or deficient decision rules. A major ultimate goal, outside the scope of our present study, is to suggest debiasing and other corrective procedures to enhance the quality of human information processing and decision rule selection.

This mixed scanning based planning and decision support system is based upon rational search for a dominance structure which will enable exposure of some of the processes upon which judgment and choice is based. In particular, it enables determination of the precise point in a dominance structure search process when a decisionmaker is able to select a single non-dominated alternative. We should be able to do this without resorting to a complete elicitation of precise parameter information and prioritization of all alternatives. The activity of complete precise determination of all parameter information is often stressful and time consuming, may require perspectives outside of the experiential familiarity of the decisionmaker, and allows few results until conclusion of the aiding effort.

The overall process described here appears well suited to accommodating the fact that neither individuals nor groups possess static decision styles capable of being stereotyped and captured by a rigid, inflexible support process. It is specifically recognized that an interactive process is needed that is capable of adaptation to a variety of decision styles that are contingency task structure dependent. System design should reflect the realization that is generally not possible to define a problem or issue fully until one knows potential solutions to the issue. A major cause of this is the fact that information to fully define the issue generally becomes available only as one evaluates potential solutions. Planning and decisionmaking will therefore necessarily be iterative.

4. Behavioral Relevance Issues

Our decision support system design paradigm is based upon a process model of decisionmaking in which a person perceives an issue which may require a change in the existing course of action. On the basis of a framing of the decision situation, one or more alternative courses of action, in addition to the present option which may be continued, are identified. A preliminary screening of the alternatives, using conjunctive and disjunctive scanning, may eliminate all but one alternative course of action. Unconflicted adherence to the present course of action or unconflicted change to a new option may well be the meta strategy for judgment and choice that is adopted if the decisionmaker perceives that the decision situation is a familiar one and that the stakes are not so high that a more thorough search and deliberation is needed [11].

Alternately, if the decision environment is an unfamiliar one, or the stakes associated with judgment and choice are high, a more vigilant form of informative acquisition, analysis, and interpretation are called for. This desire for more vigilant information processing leads to a search for a dominance pattern among alternatives, the search for new alternatives that are not dominated by presently identified alternatives, and the elimination from further consideration of alternatives that are dominated. If no single non-dominated alternative is found, adjustments to the dominance structure of alternatives are made through various forms of cognitive activity such as: attribute aggregation, additional information acquisition and analysis, and identification of additional attributes and/or alternatives. This is continued until the structure of needs, objectives, attributes, and alternative action options, and their impacts are such that identification of a single non-dominated alternative results. This "single alternative" may well represent a combination of subalternatives. If there is insufficient time and experience to accomplish these cognitive activities, hypervigilance generally results. The decisionmaker is in a situation where the present course of action is diagnosed as unfortunate and there is a shortage of time and experience that might enable identification and evaluation of an appropriate one.

Given sufficient time and experience, vigilant information processing often results from the aforementioned tasks. Figure 1 presents some salient features of this dominance process model for search, discovery, judgment, and choice.

The mode of judgment and choice that is "proper" depends upon the decisionmaker's situation diagnosis of the contingency task structure. Here, "proper" decision behavior is based upon the assumption that the environment, the task, the experiential familiarity with the task, and the environment that constitutes the contingency task structure are diagnosed correctly. If this is not the case, then the strategies leading to unconflicted change, adherence, or vigilant information processing may be significantly flawed. The role of the contingency task structure in situation diagnosis and in influencing, at a meta level or systems management level, the process of judgment and choice is, therefore, a very important one.

There have been many realistic paradigms of the process of judgment and choice. We believe that the dominance process model described here is not inconsistent with the primary features and intensions of these descriptive models. Our purpose, however, is to develop a conceptual design for a prescriptive approach to judgment and choice that will aid in the search for better decisions. We recognize that a truly rational approach to prescriptive decisionmaking must be cognizant of the process of decisionmaking as it evolves in a descriptive fashion, that is to say process rationality, or it will not be possible to evolve substantively rational support systems.

It is important that an appropriate decision support system be capable of assisting the decisionmaker through encouragement of full information acquisition, including that which may disconfirm strongly held beliefs, and in the analysis and interpretation of this information such as to avoid a variety of cognitive biases and poor information processing heuristics that may lead to flawed judgment and choice [2,12].

A realistic decision support process is necessarily iterative. Several desiderata follow from this:

1. We should allow for top-down or bottom-up structuring of the attributes of outcomes, or impacts of decisions. The "tree" or "hierarchy" of attributes should be structured to the depth believed appropriate by the decisionmaker.
2. Rather than force a decision situation structural model in the form of a tree, we should encourage the decisionmaker to identify a cognitive map of goals, objectives, needs, attributes, alternatives, and impacts that is reflective of the way in which the decisionmaker perceives diagnostic and causal inferences to occur. At some later time this cognitive map may be used to structure a multinode decision tree which is representative of substantive rationality, but not at all necessarily representative of process rationality.
3. We should encourage identification of alternative courses of action, additional attributes of decision outcomes, and revisions to previously obtained elicitations, at any point in the decision support process as awareness of the decision situation and its structure grows through use of the support system.
4. We should not force a person to quantify parameters to the extent that this becomes overly stressful, or behaviorally and physically irrelevant in view of the inherent uncertainties or imprecision that is associated with the knowledge of parameters characterizing the decision situation structural model or their assessment.

These have two primary implications with respect to our interpretation efforts. We allow for revision in the elicited structure of the decision situation and for the identification of new options as awareness of the decision situation grows. Also, we do not require the decisionmaker to quantify parameters beyond the level felt appropriate for the situation at hand. If the decisionmaker feels comfortable in exercising precision with respect to factual outcomes, this is perfectly acceptable and desirable. But parameter imprecision should be allowed if we are to have a realistic support process.

ARIADNE allows parameter imprecision in order to satisfy this quantification relevancy requirement, as do approaches based on fuzzy set theory [16]. We encourage the decisionmaker to specify precise values or numerical ranges for facts and values. Thus we allow, for example, expressions for alternative (A) scores on attributes (i) in the form $0.2 \leq v_i(A) \leq 0.5$, weights associated with attribute i in the form $0.2 \leq w_i \leq 0.4$, and probabilities of event (i) resulting from alternative A in the form $0.3 \leq p_i(A) \leq 0.45$. We allow ordinal representations in the linear forms $v_i(A) \leq v_i(B) \leq v_i(C)$, $2w_i < w_j < w_k$, $p_j(A) \leq p_i(A) \leq 3p_k(A)$, or in similar forms. Quantification of imprecision in the form of numerical bounds on parameters always leads to what we call "behaviorally consistent information sets (BCIS)." Sometimes totally ordinal information may need further quantification in order to make the precision and rigidity of the mathematical correspond to the intensions of the decisionmaker in making a purely ordinal specification. This is generally not needed to obtain solutions but, rather, to

obtain parametric models that are faithful to the understandings of the decisionmaker. For example that ordinal alternative score inequalities $0 \leq v_i(A) \leq v_i(B) \leq v_i(C) \leq 1$ are satisfied by the relations $0 \leq v_i(A) \leq 1-2t$, $W \leq v_i(B) \leq 1-t$, $2W \leq v_i(C) \leq 1$ for small positive t and W which in the limit become zero. It will generally not be the case that the decisionmaker would express this much imprecision, and would wish to see it more fully quantified to be reflective of (subjective) beliefs. It is, therefore, important that a simple and informative display of value scores, weights, and probabilities be provided to the decisionmaker. This will enhance interactive use of the support system and will enable learning of the impact of these parameters, and associated imprecision, upon decisions.

5. Conclusions

In this paper, we have examined some underlying considerations that have influenced the development of a decision support system that specifically recognizes that imprecise and incomplete knowledge is important to judgment and choice and which allows for its incorporation in the knowledge base of a decision support system. The system allows for judgment and choice at a skill based wholistic level as well as at the formal reasoning based level at which most decision analysis based paradigms operate. For detailed discussions of the algorithmic content of ARIADNE, the reader is referred to [6,7,14-26].

References

- (1) Rasmussen, J., "Skills, Rules, and Knowledge: Signals, Signs, and Symbols, and Other Distinctions in Human Performance Models," IEEE Transactions on Systems Man and Cybernetics, Vol. 13, No. 3, May/June 1983.
- (2) Sage, A. P., "Behavioral and Organizational Considerations in the Design of Information Systems and Processes for Planning and Decision Support," IEEE Transactions on Systems Man and Cybernetics, Vol. SMC-11, No. 9, September 1981, pp. 640-678.
- (3) Pearl, J., Leal, A., and Saleh, J., "GODDESS: A Goal Directed Decision Structuring System," IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. PAMI-4, No. 3, May 1982, pp. 250-262.
- (4) Rajala, D. W., and Sage, A. P., "On Decision Situation Structural Modeling," Policy Analysis and Information Science, Vol. 4, No. 1, July 1980, pp. 53-81.
- (5) Sage, A. P., "On Sensitivity Analysis in Systems for Planning and Decision Support," Journal of the Franklin Institute, Vol. 312, No. 3/4, September/October 1981, pp. 265-291.
- (6) White, C. C., Sage, A. P., and Dozono, S., "Implications of Imprecise Measure of Preference on Alternative Ranking Specificity," Proceedings of 6th IFAC Symposium on Identification and System Parameter Estimation, Washington, DC, June 1982, pp. 592-597.
- (7) White, C. C., Sage, A. P., and Dozono, S., "A Model of Multiattribute Decisionmaking and Trade-off Weight Determination Under Uncertainty," IEEE Transactions on Systems, Man and Cybernetics, Vol. 14, No. 1 January/February 1984.
- (8) Hammond, K. R., McClelland, G. H., and Mumpower, J., Human Judgment and Decision Making, Praeger, 1980.
- (9) Argyris, C., Reasoning, Learning, and Action, Jossey-Bass, 1982.
- (10) Sage, A. P., Methodology for Large Scale Systems, McGraw-Hill Book Co., 1977.
- (11) Janis, I. L., and Mann, L., Decision Making, Free Press, 1977.
- (12) Kahneman, D., Slovic, P., and Tversky, A. (eds.), Judgment Under Uncertainty: Heuristics and Biases, Oxford University Press, 1982.
- (13) Fischhoff, B., Slovic, P., and Lichtenstein, S., "Knowing What you Want: Measuring Labile Values," in T. S. Wallsten (ed.), Cognitive Processes in Choice and Decision Behavior, Erlbaum, 1980.
- (14) White, C. C. and Sage, A. P., "Second Order Stochastic Dominance for Multiple Criteria Evaluation and Decision Support," Proceedings 1981 International Conference on Cybernetics and Society, Atlanta, GA, October 1981, pp. 572-567.
- (15) White, C. C. and Sage, A. P., "Multiple Objective Evaluation and Choicemaking Under Risk with Partial Preference Information," International Journal on Systems Science, Vol. 14, No. 5, 1983, pp. 467-485.
- (16) White, C. C., Sage, A. P., and Scherer, W. T., "Decision Support with Partially Identified Parameters," Large Scale Systems, Vol. 3, No. 3, August 1982, pp. 177-190.
- (17) White, C. C., Lozono, S., and Scherer, W. T., "An Interactive Procedure for Aiding Multiattribute Alternative Selection," Omega, Vol. 11, No. 2, 1982, pp. 212-214.
- (18) White, C. C. and Sage, A. P., "Imprecise Importance Weight Assessment for Multilevel Objectives Hierarchies," Proceedings IEEE International Conference on Large Scale Systems, Virginia Beach, VA, October 1982, pp. 298-301.
- (19) Sage, A. P. and White, C. C., "A Top Down Approach to Imprecise Interactive Attribute Weight, Structure, and Alternative Score Assessment," Proceedings IEEE International Conference on Systems Man and Cybernetics, Seattle, WA, October 1982, pp. 663-667.
- (20) Sage, A. P., and White, C. C., "Behavioral Issues in a Knowledge Based Interactive System for Decision Support," Proceedings IEEE Conference on Decision and Control, December 1982, pp. 1049-1052.
- (21) White, C. C. and El Deib, H., "Multistage Decisionmaking with Imprecise Utilities," in Essays and Surveys on Multiple Criteria Decision Making, Springer-Verlag, New York, 1983, pp. 400-405.
- (22) White, C. C., "Sequential Decisionmaking with Uncertain Future Preferences," Operations Research, to be published.
- (23) White, C. C. and El Deib, H., "Parameter Imprecision in Finite State, Finite Action Dynamic Programs," submitted for publication.

- (24) White, C. C., and Sage, A. P., Dozono, S., and Scherer, W. T., "An Evaluation of ARIADNE," Proceedings Sixth MIT/ONR Workshop on Command and Control, July 1983.
- (25) Sage, A. P. and White, C. C., "ARIADNE: A Knowledge Based Interactive System on Planning and Decision Support," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC 14, No. 1, January/February 1984.
- (26) White, C. C. and Sage, A. P., "A Multiple Objective Optimization Based Approach to Choice-making," IEEE Transactions on Systems Man and Cybernetics, Vol. SMC-10, No. 6, pp. 315-326, June 1980, reprinted as a chapter in R. Sacks and J. D. Palmer (eds.) The World of Large Scale Systems, IEEE Press and John Wiley and Sons, Inc., 1982, pp. 335-346.

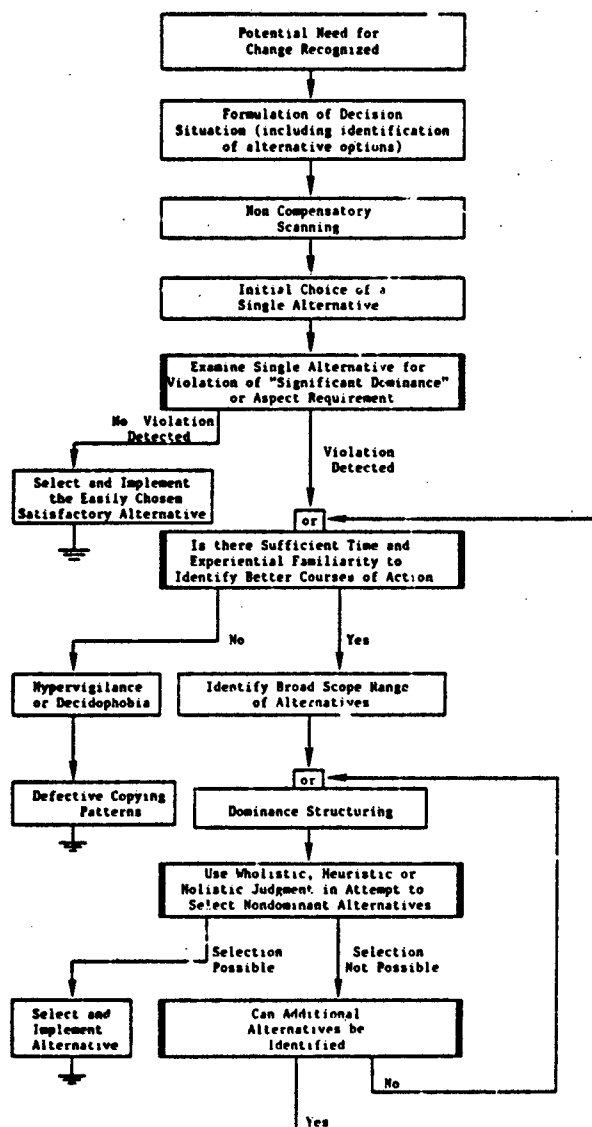


Figure 1. Descriptive Dominance Structural Model of Decision Process.

AN EVALUATION OF ARIADNE

by

Chelsea C. White, III, Andrew P. Sage, Shigeru Dozono,
and William T. Schirer

Department of Systems Engineering
University of Virginia
Charlottesville, VA 22901

ABSTRACT

In this paper, we present the objectives, operational details, results, and conclusions of an evaluation of a decision aiding procedure ARIADNE. The results of the evaluation indicate that ARIADNE, in comparison to a well-known decision aiding procedure called SMART, (1) has a more flexible model of parameter value description that tends to reduce assessment stress and makes ARIADNE more useful in situations where information precision is poor, (2) allows earlier presentation of initial alternative ranking information and (3) allows the decisionmaker to adjust the mix of alternative ranking specificity and parameter value precision.

I. INTRODUCTION

In this paper, we present the objectives, operational details, results, and conclusions of an evaluation of two decision aiding approaches, ARIADNE and SMART. An indepth description of SMART (Simple, MultiAttribute Rating Technique) can be found in Edwards (1977); Detailed discussions and algorithmic descriptions of ARIADNE (Alternative Ranking Interactive Aid based on Dominance structure information Elicitation) are presented in the companion paper (Sage and White, 1983). In Section II, we list the hypotheses that were tested during the evaluation. The operational aspects of the evaluation are detailed in Section III. We present the results of the evaluation, and the conclusions that we obtain from these relative to the identified hypotheses, in Section IV.

II. IDENTIFIED HYPOTHESES:

The intent of the evaluation was to test the following hypotheses:

1. Use of the more general model of parameter value description in ARIADNE tends to reduce stress associated with the assessment of alternative values and attribute weights.
2. Use of the more general model of parameter value description in ARIADNE tends to increase confidence in the final alternative selected.
3. The ability to provide additional parameter information, in a form and sequence selected by the decisionmaker, and to observe its impact on alternative ranking in an iterative fashion is a desirable feature of ARIADNE.

4. ARIADNE requires less time for use than does SMART.
5. Decisionmakers do not feel that it is necessary for an aid to produce a single best alternative to assist the decisionmaker in selecting the most preferred alternative.
6. ARIADNE is more useful than SMART in situations where information precision is poor.
7. Problems typically encountered in the subjects' operational environment would be more appropriately examined aided by ARIADNE than aided by SMART.
8. ARIADNE is no more difficult to understand and use than SMART.

We now discuss the procedures for testing these hypotheses.

III. OPERATIONAL ASPECTS OF EVALUATION

Eight (8) civilians employed by the United States Army Foreign Science and Technology Center (FSTC) participated as subjects in the evaluation. Each of the subjects had had extensive involvement in technical project evaluation in a military environment and thus had sufficient experience to appreciate the difficulties and operational issues involving proposal evaluation. Proposal evaluation was a subject addressed in the more specialized of the two decisionmaking scenarios examined by the subjects during the evaluation.

The two scenarios developed for the evaluation were proposal evaluation and sports car selection. Each of these scenarios can be obtained from the authors. It was assumed that each scenario involved decisionmaking under certainty. The first scenario was designed to represent a realistic proposal evaluation problem that might occur in a DOD funding agency. Although attributes and some ordinal relations among attribute weights were specified in the RFP to which the proposals were to respond, information presented in the five (5) submitted proposals from which to deduce utility scores and hence tradeoff weights was often vague and/or not available. Also, there was room for judgement in strengthening the ordinal relations among the attribute weights that were provided in the RFP summarized in the scenario. The sports car selection scenario was designed to represent a much more precisely defined alternative selection problem.

Standard procedures were used in order to investigate and compensate for effects due to facilitation style and order with respect to decision aid and scenario.

* This research has been supported by the Office of Naval Research under contract number N00014-80-C-0542.

The chronology for the evaluation was as follows:

1. General briefing. A briefing was given to the volunteer subjects regarding the purpose of the evaluation and the characteristics of the two aids. Individual evaluation sessions were scheduled, and both scenarios were given to the subjects to read prior to the individual sessions.

2. Individual sessions. Individual evaluation sessions were conducted. Each session for each individual subject involved a subject, facilitator, and computer terminal operator. If ARIADNE was used, assessed information regarding lowest level attribute utility scores and tradeoff weights was allowed to be imprecise and was translated into linear inequalities by the facilitator and/or computer terminal operator. Initially, only utility score information for the alternatives on the identified attributes was assessed. Once this assessment was completed, a domination digraph on the alternatives was computed and displayed to the subject. The subject could also view a score sheet of values from which this digraph was obtained. If this digraph provided sufficient information for alternative selection, then this portion of the session was halted. If not, then further utility score and/or tradeoff weight information was requested and the resulting domination digraph displayed. This information could concern attribute scores and weights not previously obtained or more precise estimates of previously elicited scores and weights. This iterative procedure continued until the subject halted the process.

If SMART was used, all parameters were precisely assessed. Then the total linear order on the alternatives was displayed. If the subject wished, a post optimal sensitivity analysis was performed on whatever single parameter values were of concern. Detailed descriptions of facilitation protocols can be obtained from the authors.

During the examinations of each scenario, the computer terminal operator completed an Analyst Information Sheet, detailing various times and types of requests. After completing both scenarios, the subjects were asked to complete a short questionnaire and return it. Copies of both the Analyst Information Sheet and the Questionnaire can be obtained from the authors.


IV. RESULTS AND CONCLUSIONS:

We now examine each of the hypotheses in the context of data collected during the evaluation.

Hypothesis # 1: Use of the more general model of parameter value description in ARIADNE tends to reduce stress associated with the assessment of alternative values and attribute weights. Relevant Questionnaire questions (Q) and responses (R):

Q1: Being allowed to express parameter values imprecisely using ARIADNE produced more stress than being required to state all parameter values precisely using SMART.

R: (presenting individual subject scores below the line; +5(-5) indicates strong (dis)agreement)



indicating a tendency to agree with the hypothesis. There were two relevant comments: S1 (Subject 1): "liked the flexibility."


S3: "Perhaps more stress initially (with ARIADNE) because I didn't understand the process. But after using it, I would be inclined to say far less stress (with ARIADNE) for complex applications."

Q6: What was the most comfortable way of expressing parameter value information for you?

R: Indicated that two subjects (S2, S3) preferred exact values, two subjects (S1, S8) preferred interval estimates, and four (S3, S4, S6, S7) preferred ranking statements for expressing parameter value information. There was one relevant comment from S3: "for the exercise today, [exact value] would be the answer; however, for actual application in complex areas, [ranking statements] is my answer." These responses cause us to conjecture that if the subjects had been more experienced in expressing parameter values imprecisely (several had experience in expressing value scores and tradeoff weights precisely) then there would have been stronger support for this hypothesis.

Hypothesis # 2: Use of the more general model of parameter value description in ARIADNE tends to increase confidence in the final alternative selected. Relevant Questionnaire questions and responses:

Q2: I felt more confidence in the final alternatives chosen when aided by ARIADNE than in the final alternative produced when aided by SMART.



indicating slight support for the hypothesis. There were no especially relevant written comments.

Q8: I disagreed with the action alternative recommendation obtained using ARIADNE.

R: (S1,...,S7) "no"; S8 "don't remember."

Q14: I disagreed with the action alternative recommendation obtained using SMART.

R: S2 "yes;" (S1, S3, ..., S7) "no;" S8 "don't remember." Thus, responses to questions the latter 2 indicate no recognizable difference in the perceived quality of the decisions made using ARIADNE versus those made using SMART. There were no relevant comments associated with either question.

Q17: Which approach would you prefer to use to make recommendations to others concerning evaluation and prioritization and why?

R: (S1, S3, S7, S8) "ARIADNE;" (S2, S5, S6) "SMART;" S4 "neither." Three relevant comments were made: S4: "Depends on Situation." S6: "Since SMART gives percentages, one can see if two proposals are close and then adjust the rankings by consideration of factors that were not originally considered." S7: "By only indicating preference as opposed to exact values is very helpful; weight max-min display nice feature."



We feel S4's comments tends to explain his response to this question. The comments of S6 and S7 indicate that the max-min display (indicating the maximum and minimum values of expected utility for each alternative)

incorporated into ARIADNE was liked by S7 and may not have been requested by S6. We conjecture that had S6 seen this display, his response would have been different.

Q12: Use of ARIADNE encouraged me to carefully weigh the positive and negative consequences of each alternative.

Q16: Use of SMART encouraged me to carefully weigh the positive and negative consequences of each alternative.

Responses to Q12 and Q16 were, respectively:




indicating that ARIADNE and SMART provided an approximately equal level of encouragement to the subject for him to carefully weigh the positive and negative consequences of each alternative.

In summary, questionnaire responses indicate that the level of confidence in the output of ARIADNE and SMART appear to be quite similar. Also indicated was that both aids equally encourage the careful weighing of the possible consequences of the alternatives.

Hypothesis # 3. The ability to provide additional parameter information, in a form and sequence selected by the decisionmaker, and to observe its impact on alternative ranking in an iterative fashion is a desirable feature of ARIADNE. Relevant Questionnaire questions and responses:

Q3: Being able to provide additional parameter information and then to observe its impact on alternative ranking in an iterative fashion, was a desirable feature of ARIADNE.




indicating strong support for the hypothesis.

As an indication of how often the feature of ARIADNE being evaluated in this hypothesis test was exercised, the computer terminal operator recorded on the analyst information sheet the number of iterations required in constructing the final domination digraph. The number of iterations for each subject was respectively: 3, 3, 0, 4, 0, 2, 5, and 3.

Thus the subjects developed a reasonable level of experience with this feature of ARIADNE, placing a high level of confidence in the responses to Q3.

Q4: Knowing that a single iteration always produces a best alternative is a desirable feature of SMART.




indicating that the subjects found moderately attractive the fact that SMART always produces a best alternative in one iteration.

The responses to Q3 and Q4 indicate that the decisionmaker should be encouraged to be as precise as possible in order to reduce the number of iterations of ARIADNE for final decision selection.

Hypothesis # 4: ARIADNE requires less time for use than does SMART. Relevant questions and responses:

Q5: Use of SMART lead me to a decision more quickly than use of ARIADNE.



indicating that the perceptions of the subjects tend not to support the hypothesis.

Timing data recorded on the analyst information sheets support the perceptions of the subjects with regard to the average total length of time per session. Let VET = value elicitation time, WET = weight elicitation time, TST = total session time. Then, the timing data (in minutes) are as follows:

	1	2	3	4
	APS			
VET	41.75	18.75	19.25	21.25
WET	9.5	10.5	6.0	14.0
TST	57.0	41.25	47.25	51.0
	CCW			
VET	29.5	24.25	28.75	45.0
WET	11.75	9.75	10.0	11.5
TST	57.5	40.75	47.0	51.25
	TOTAL			
VET	35.625	21.5	24.0	23.125
WET	10.625	10.125	8.00	12.75
TST	57.25	41.0	47.125	51.125

where 1 = ARIADNE, 2 = SMART, 3 = proposal evaluation, 4 = sports car selection. Thus, use of ARIADNE required on average 40% more total time than did SMART and the proposal evaluation scenario required on average 8.5% more time to evaluate than did the sports car selection scenario. Both facilitators on average required 49.125 minutes per scenario.

Also recorded on the analyst information sheet was the length of time between the beginning of the session and the beginning of the weight elicitation process. For ARIADNE, this length of time represents an upper bound on the length of time to the presentation of the first digraph. These times for the 8 subjects were: 25, 41, 20, 30, 30, 15, 25, 35 for an average of 27.625 seconds. We remark that SMART required approximately 48% more time to provide initial alternative ranking information to the decisionmaker than did ARIADNE. If we assume that total session times for SMART and the lengths of time between the beginning of the session and the beginning of the weight elicitation process for ARIADNE are realizations of normally distributed random variables with unknown means and unknown variances, then a standard statistical test indicates that these realizations come from two different random variables with a confidence level of greater than 0.95.

SMART) would cause ARIADNE to be perceived by the user as no more difficult to understand and use than SMART. Training time to achieve such a level of familiarity is likely to be longer for ARIADNE than for SMART due to the relative increased flexibility inherent in ARIADNE.

Other issues and associated questions:

Q18: Was the posterior sensitivity analysis associated with SMART helpful?

Q19: Which decision making scenario was the most appropriate for ARIADNE and why?

Q20: Which decision making scenario was the most appropriate for SMART and why?

Responses to Q18, Q19, and Q20 indicated, respectively, that: 1. the post-optimality sensitivity analysis feature was useful in SMART, 2. the proposal evaluation scenario was most appropriately evaluated using ARIADNE, and 3. the sport car selection scenario was most appropriately evaluated using SMART. Generally, comments to questions # 19 and # 20 indicated that for complex decision selection situations having less quantitative information available, ARIADNE would be preferred to SMART, which provides further support for hypothesis # 6.

Summary, our evaluation lends credence to the following claims:

1. The more flexible model of parameter value description employed by ARIADNE tends to reduce assessment stress and makes ARIADNE more useful than SMART in situations where information precision is poor.
2. The iterative, progressive information requirements associated with ARIADNE is a desirable feature that allows earlier presentation of initial alternative ranking information than does SMART.
3. SMART requires less total time for use than does ARIADNE.
4. Being able to adjust the mix of alternative ranking specificity and parameter value precision is a desirable feature of ARIADNE.
5. ARIADNE may require more training than would be required by SMART for successful use.
6. The level of confidence in the output of ARIADNE and SMART appear to be quite similar.

REFERENCES

1. Edwards, W., "How to Use Multiattribute Utility Measurement for Social Decisionmaking," IEEE Systems, Man, and Cybernetics, Vol. SMC-7, pp. 326-339, 1977.
2. Sage, A.P., and White, C.C., "A Knowledge-Based, Interactive Procedure for Planning and Decision Support Under Uncertainty and Parameter Imprecision," 6th MIT/ONR Workshop, 1983.
3. White, C.C., Dozono, S., and Scherer, W.T., "An Interactive Procedure for Aiding Multiattribute Alternative Selection," Omega, Vol. 11, 1983.

4. White, C.C., Wilson, E.C., and Weaver, A.G., "Decision Aid Development for Use in Ambulatory Health Care Settings," Operations Research, Vol. 30, 1982, pp. 446-463.
5. Sage, A. P., and White, C. C., "ARIADNE: A Knowledge Based Interactive System for Planning and Decision Support," IEEE Transactions on Systems, Man and Cybernetics, Vol. 14, No. 1, January/February 1984.
6. Sage, A. P., "A Methodological Framework for Systemic Design and Evaluation of Planning and Decision Support Systems," Computers and Electrical Engineering, Vol. 8, No. 2, June 1981, pp. 87-102.

RESEARCH ON COGNITIVE COLLABORATION BETWEEN PERSONS AND COMPUTERS

Marvin S. Cohen

Decision Science Consortium, Inc.
7700 Leesburg Pike, Suite 421
Falls Church, Virginia 22043

Abstract

The introduction of decision aids and knowledge-based expert systems incurs resistance when non-congenial styles of problem solving are imposed on users. On-going research addresses the design of computer-based display and analysis systems which cater flexibly to personal styles while providing non-obtrusive safeguards against potential errors and biases. Capabilities which permit monitoring of the user's task by the computer and of the computer by the user have been explored.

The Problem

High-level users of computer-based information systems typically find that either too little or too much help is offered.[1][2] On the one hand, sophisticated systems are available for data retrieval, analysis, and display, yet they provide little guidance in selecting the information that ought to be retrieved or the type of analysis which the user ought to apply. On the other hand, decision aids and knowledge-based expert systems typically impose an analytical structure and mode of interaction which may prove inappropriate or uncongenial to the user's own preferred style of problem solving. Users, in short, are caught between systems that automate routine functions and systems which cannot help but dominate any dialogue with the decision maker.

It might be thought that as computer-based systems more completely automate intellectual tasks, the issue of user preferences will become moot. Yet the most critical characteristic of these new applications is that they are neither fully objective nor demonstrably optimal. Knowledge-based expert systems incorporate the assumptions and modes of reasoning employed by human specialists. Decision-analytic aids provide logical constraints for inputs from human experts or decision makers regarding subjective probabilities, preferences, and problem structure. Both kinds of systems are appropriately regarded only as fallible advisors. Complete automation could be inappropriate if users possess substantive expertise or analytic insights not incorporated in the computer.

What is required, both to encourage user acceptance and to enhance aid performance, is a repertoire of techniques for blending the expertise of the user and computer. Such techniques must be fine-grained and flexible enough to capture shifting availabilities of human and computer resources, relative levels of expertise, and user preferences.

Unfortunately, in the design of systems that foster cognitive collaboration, two basic objectives tend to conflict: On the one hand, we want to exploit user inputs where (and only where) they can enhance the overall credibility of aid outputs. On the other hand, users have their own preferences and styles of problem solving that may not correspond to optimal patterns of

allocating cognitive effort. By imposing a rigid structure on person-machine interaction (however "optimal" it may be from the point of view of relative expertise), the net outcome may be less effective problem-solving--including perhaps a failure to use the system altogether.

To deal with these conflicting objectives, our research has focused on three broad capabilities in cognitive system design:

- flexible blending of computer and human contributions, under the personal control of the user;
- monitoring by the computer of selected human-performed tasks; and
- monitoring by the human of selected computer-performed tasks.

The first principle maximizes the tailoring of person-computer interaction to the particular style of a user. The second and third principles provide a prescriptive counterbalance: they are designed to compensate for deviations from optimality that may emerge from the first principle, and to do so in the most non-obtrusive way possible.

In the following sections, we briefly summarize some of the research we have done under these three headings. The focus is on the psychological underpinnings and implications of the work, rather than on the details of the decision aids that have been developed. This work has been supported by the Engineering Psychology Group of the Office of Naval Research under two on-going contracts.*

Aids for Personalized Decision Making

Under the Defense Department's Small-Business Advanced Technology (DJSAT) program, DSC has explored the design of a computer-based display and analysis system which is customized to the personal cognitive styles of users.[3] The design process has drawn on relevant work in the cognitive psychology of judgment and choice, in computer science, and in the prescriptive theory of decision making. A prototype system, developed for attack submarine antisubmarine warfare (ASW), is based in part on our own study of individual differences in decision-making styles among submarine officers.

The Decision Setting and the Decision Process

The dilemmas faced by the command staff of a hunter-killer submarine in approaching and attacking an (as yet) unalerted hostile submarine are characteristic of situations involving stealth in warfare: How long should I attempt to remain undetected and to improve my position, before I tip my hand by launching a weapon? In planning an attack, the Commander faces a number of choices (among weapons, targets, approach maneuvers, and times of fire) and is flooded with an increasingly

manageable quantity of data (about the target, own ship, and environment). To capitalize on the element of surprise, a price must be paid in the quality of the data, the complexity of options, and the strenuousness of the choice process. In all of these areas, there is substantial leeway for differences in individual cognitive styles of coping.

Situation Assessment. Assessment tasks must depend almost exclusively on passive sensors (which do not alert the enemy); as a result, data are often fragmentary, noisy, and inconsistent. Little or no guidance is provided in reconciling multiple conflicting estimates of the same variable (e.g., target range), organizing data acquisition, assessing the quality of estimates, or drawing inferences about critical opportunities and dangers (e.g., probability of kill, probability of counterdetection).

Work in cognitive psychology suggests a number of ways in which people may simplify the cognitive demands of these tasks at the risk of suboptimal performance. Where multiple estimates are available for a single variable (e.g., target range), people tend to ignore evidence that contradicts a favored, or earlier, datum and to double count redundant evidence.[4] Patterns of information search tend to avoid stringent tests of favored hypotheses.[5][6][7] Assessments of degree of certainty tend to be overconfident.[8] When inference proceeds in stages (e.g., deriving probability of kill from information about range, which is derived from bearings data), people often act as if conclusions at earlier stages were known to be true, rather than merely inferred.[9] Similarly, the probability of a detailed scenario is often judged higher than the probabilities for component events.[10]

Option Generation. Interdependent elements of a tactic should be considered together: for example, use of certain types of weapons may be precluded by the risk of counterdetection by a third party threat, unless appropriate maneuvers, firing position, and time of fire are selected. The consequences of immediate decisions for later choices may also be critical - e.g., the ability to proceed against or evade a second threat after the initial attack, or the ability to respond if unexpectedly counterdetected.

Research suggests that the process of formulating options is often truncated in a variety of ways. People prefer to treat the elements of complex options as if they were independent choices. There is a tendency to formulate options that span only a short time-frame, and to overlook, as a result, the cumulative risk of pursuing a given course of action over a long period of time.[11] Individuals differ in the degree to which future acts are considered in current planning [12] and in the sheer number of options they consider.[13] Customary ways of viewing a problem tend to hinder the generation of novel and creative solutions.[14]

Choice. The aim of avoiding counterdetection frequently clashes with other goals. A premature attack may both alert the enemy and miss; yet continued approach increases the risk that own ship will be detected before attack or that the target will successfully evade. Perhaps because the information load tends to be large, simple heuristic decision rules are often invoked: e.g., for time-of-fire, "avoid counterdetection"; or "fire as soon as within maximum weapon range and in possession of a range solution."

There is a growing literature in cognitive psychology suggesting that rules like these may be adopted to reduce the cognitive effort that would be involved in a thorough consideration of each option.[15][16] With one

such rule, Elimination-by-Aspects [17], all options falling below a cutoff on an attribute are eliminated, and attributes are considered in turn until only one option remains. In "satisficing" [18][19], the decision maker considers a sequence of options but stops as soon as he finds one that clears a cutoff or set of cutoffs on selected attributes. In each of these examples, an option might be eliminated even though it scores very high on some dimensions. In the submarine context, such rules exclude a balancing of tradeoffs - such as accepting a small risk of detection in order to accomplish a mission objective. Satisficing can cause superior options (e.g., a later time of fire) to be overlooked.

It has been suggested that experts differ from novices in their capability to individually recognize a very large number of different problem situations. [20] Klein [21] argues that experts tend to reason holistically, by analogy with previous similar experiences, rather than by analysis and computation. To the extent that this is true, we might expect that for experienced commanders all stages of decision making - situation assessment, option generation, and choice - would be considerably streamlined. At the least, we would expect decision makers to differ in the degree to which they arrive at highly integrative conclusions without the necessity of a large number of explicit intervening steps.

Individual Differences in Decision-Making Style

Early in the design process of the prototype aid, data regarding individual patterns in the use of information was gathered in a procedure involving four former submarine command personnel. They received a questionnaire describing a realistic multiple threat ASW approach and attack scenario. The questions were designed to focus not only on observable patterns of information use, but also on the less conspicuous decision-making processes within which that information plays a role. At each of a number of break points in the scenarios, the officers were asked to specify: the information currently available on board the submarine which they would seek, the source from which they would seek it, the combat decisions that depended on the information, the way the information would affect those decisions, and the objectives of the decision.

Analysis of this data suggested that there were important differences in styles of data gathering, option formulation, and choice to which an aid might cater.

Situation Assessment (A): Amount of Information. The total number of items utilized varied considerably, from 42 information requests by one officer to 18 by another.

Situation Assessment (B): Information Search Pattern. Requests for data fell into two quite distinct patterns. Two of the officers tended to organize data acquisition by source, asking for a "dump" of current estimates from sonar, plot, or fire control, then going on to another source. The other two officers organized data acquisition by item, asking for a given estimate, like target range, from a variety of sources or else selectively requesting different items from different sources.

Option Generation: Time-Span. The officers differed in the time horizon of the options they considered, e.g., focusing exclusively on the immediate actions required to regain a lost contact versus evaluating in advance approach tactics contingent upon recovery of the contact.

Choice (A): Level of Integration. There were differences among officers and for all individual officers across situations in the scope of the objectives which they brought to bear on the evaluation of options. Objectives might be specified quite broadly as preserving own ship, or more narrowly as avoiding counterdetection or watching for clues regarding counterdetection status. Similarly, the goal might be killing the target, achieving a suitable firing position, or opening torpedo tube doors.

Choice (B): Number of Evaluative Dimensions. One officer combined concerns for own ship survival and killing the target in all decisions (each concern might be at various levels of integration). Two of the officers appeared to shift back and forth in their focus between these concerns. The fourth officer went all out for target kill, never once explicitly mentioning an objective related to own ship survival (at any level of integration).

Choice (C): Use of Cutoff Criteria. Three of the four officers evaluated actions explicitly in terms of cutoffs. All three used the achievement of maximum weapon range as a criterion for attack; one used arrival at counterdetection range as a criterion for withdrawal.

A Prototype Personalized System

A prototype personalized aid has been designed and partially implemented for approach and attack planning by the command staff of a nuclear attack submarine. However, only the data base of the aid is affected by the nature of the specific application. Its functional logic, and the methods used to achieve both personalization and prescriptive impact, are quite general. The implementation of a demonstration prototype system in a specific context, however, permits a realistic test of the feasibility of the concepts, with potential users.

Figure 1 outlines the general logic of the cognitive interface. The prototype aid design consists of a data base, a flexible general-purpose Planning Module, and four relatively specialized routines for customizing the aid. The system utilizes principles of spatial data management which combine an undemanding style of interaction with a high degree of user control over display contents. All user inputs are via a single simple locator device (a joystick plus button) with control properties that shift appropriately with the display region where the cursor is located.

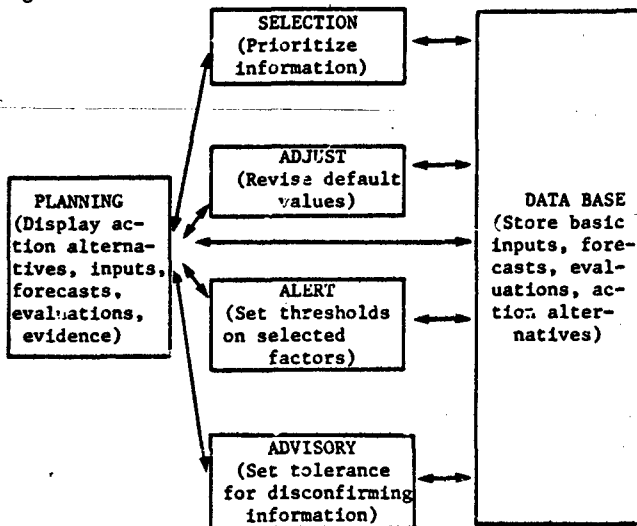


Figure 1. Structure of Prototype Aid

The display area of the Planning Module (Figure 2) is divided into a set of windows which permit simultaneous viewing of substantive results (evaluations of alternative tactics) and a variety of menus by means of which the user can specify the tactics to be evaluated, the criteria to be employed in the evaluation, and sources of validation for displayed results. A final menu enables the user to select other specialized modules (Select, Adjust, Alert, Advisory). The Planning Module facilitates a variety of personal preferences in the approach to situation assessment, formulation of options, and choice.

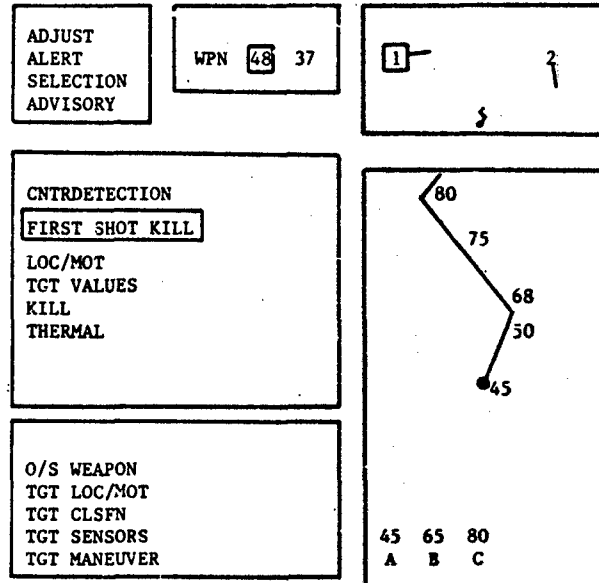


Figure 2. Planning Module Display

Situation Assessment. The data base consists of basic inputs (in the submarine testbed these concern own ship, contacts, and the environment) together with a set of prescriptive models which aggregate those inputs into high-level inferences and forecasts of critical events (e.g., counterdetection and first-shot kill) and evaluations in terms of ultimate combat objectives. The Planning Module enables users to sample information at any preferred level of aggregation in the data base. When higher-level inferences are displayed, the Planning Module clearly distinguishes conclusions from evidence, and indicates the sources from which each inference is derived. The user may elect to examine in more detail any of the evidence utilized in deriving a particular conclusion.

The Selection Module allows the user to view a map of the total data base and to personally select the portion which will be immediately accessible through the Planning Module.

The Adjust Module enables the user to insert subjective judgments in place of default values at any level in the data base. The Planning Module will then display the implications of the hypothetical or revised values for any higher-level inference. (Default values, however, continue to be stored and displayed.) The Adjust Module thus accommodates individual differences in beliefs and preferences and - from a prescriptive point of view - adds a potentially valuable source of information (the user) to the data base. We return to this feature in the last section.

The Alert Module performs situation monitoring for the user. It enables him to set a cutoff or threshold

for any variable in the data base (at any level of aggregation) when cutoffs are crossed.

Option Generation. The Planning Module facilitates the formulation and evaluation of complete tactical options (weapons, targets, approach maneuvers, and times of fire). It enables the user to vary the number of alternatives examined and the time into the future over which an option extends. A version of the aid currently under development gives the user a choice between entering his own options directly for evaluation or specifying personalized parameters to constrain automatic option generation.

Choice. In the Planning Module, the user can evaluate options by reference to objectives at any of a variety of levels of integrative scope (e.g., how quickly will the option get me to point x? How will it help improve probability of kill? What is its overall merit, combining probability of kill and probability of own ship survival?).

The Alert Module facilitates individual heuristic strategies (such as Elimination-by-Aspects and satisficing) which evaluate actions by reference to cutoffs as opposed to tradeoffs. After the user sets a threshold on a variable, the Planning Module forecasts whether or not the threshold is expected to be crossed for any action alternative which he wishes to evaluate, and if so, when. Different heuristic strategies for choice imply differences in the way information is searched: e.g., by action (run through all relevant evaluative variables for a given tactic, as in prescriptive theory or satisficing) or by criterion (examine all options for a given evaluative variable, as in Elimination-by-Aspects). In the Planning Module both of these search modes are specifically facilitated.

Prescriptive Prompting

An important factor in designing a personalized and prescriptive aid is the impact of individual preferences on outcomes. Simplifying for illustrative purposes, alternative strategies for performing the same task may fall into one of two classes in this respect:

- Strategy A is generally expected to be more accurate or yield more preferred outcomes than strategy B, but requires more training, more time, and/or draws away more attention from other tasks.

An example, in the area of choice, might be evaluating each option by reference to all the relevant dimensions (A) versus eliminating some options by reference to only a few (B). (Or, in inference tasks, ignoring important sources of uncertainty.) In these cases, differences among people in preference between A and B might reflect differences in their underlying ability to perform A, in their training or knowledge, in their handling of workload, degree of motivation, or their evaluation of the cost of errors.

- For some people, strategy A is expected to be more effective (better in accuracy, payoffs, speed, effort, etc.) than strategy B, while for other people, strategy B is more effective than A.

Payne [15] speculates that search organized by options versus search organized by attributes may reflect individual differences in the way knowledge is internally represented. People who differ in their degree of experience or areas of expertise may prefer and benefit from different ways of structuring a problem.

These distinctions have implications for the appropriateness of prescriptive advice in a personalized decision aid. In the second case discussed above, the user usually does best with the strategy which he prefers; accordingly, an interactive system should simply facilitate selection by the user of the information processing rule or structure to be employed.

In the first case, the computer's role may, at the request of the user, be somewhat more active. It involves an apparent conflict between the user-preferred and the normative strategy - though the use of the former may in fact be well justified by savings in time and effort. In such cases, the computer can assist by applying a prescriptive model to the problem, in parallel with the user's own effort which it monitors. The aid may then advise the user when discrepancies seem significant. The prescriptive model applied by the aid, of course, has no automatic claim to truth; it takes the role, rather, of a "cooperative adversary" or "devil's advocate." It enables the user to concentrate his own attention selectively, in areas that he regards as critical, while notifying him when other issues seem worthy of attention. Advisory prompts thus complement the freedom of individual choice granted by personalizing features; they encourage flexibility by offering some insurance against possible pitfalls.

Two important features of advisory prompts as we seen them are worth stressing:

- The objective is not simply to alert the user whenever there is some difference, however small, between his judgment and the output of a prescriptive model. The difference must be large enough to matter, in the actions to be selected and in their expected outcomes.[22]
- The user himself determines the size of the discrepancy that would justify a prescriptive prompt. The frequency of prompting will thus depend on his own informal assessment of the value of his time and effort relative to the cost of errors. The Adjust Module of the personalized aid enables the user to input that judgment.

Prompts may be introduced to assist users in tasks of situation assessment, option generation, and choice. Our current research involves the conceptual design, implementation, and testing of a variety of such prompts.

Situation Assessment. The user might be notified when two information sources, both of which are regarded as credible, have contradicted one another. He might then choose to implement prescriptive procedures for appropriately readjusting one or both credibility assessments downward. A prescriptive prompt might notify him on future occasions when either of the (partially) discredited sources is involved in an important conclusion.

Advisory prompts might signal when favored information search patterns seem inefficient, e.g., seeking additional confirming evidence for an already well-supported hypothesis.

Prescriptive prompts might warn users, when they estimate or subjectively adjust higher-level inferred variables, that a number of stages of uncertainty must be kept in mind. The same type of caution might be appropriate when the likelihood of a compound, or conjunctive, event is being assessed.

Option Generation. Short range planning might be

more appropriate in some situations (e.g., where feedback is continuous and mistakes can be easily and quickly corrected), while long range planning would be more suitable in others (e.g., where a risk appears small unless it is considered cumulatively over the long run). Prompts might recommend that the user consider a shift in time horizon under appropriate circumstances.

A variety of prompts might be utilized to stimulate "creativity," or the generation of novel options. The system might encourage the user to adopt, hypothetically, a new "schema" of the situation by questioning his basic assumptions about the threat, own ship, and environment - especially where the system data base actually has information that deviates from "normal" conditions. Alternatively, the system might encourage the user to better delineate the space of options by generating options tailored to single objectives, especially objectives not so far considered by the user.

Choice. Advisory prompts might signal a user who is employing cutoffs when tradeoffs bear looking into; in particular, where tradeoffs involve evaluative dimensions he has not as yet examined. More generally, the Planning Module might monitor the user's selection of information and specification of options, and derive hypotheses regarding the user's decision process and conclusions. The user would be advised when information about tactical options which he has not requested may have implications for choice that clash with the system-inferred user model.

User Override

In a personalized decision aid, ultimate control over task assignments belongs to the user. We have just seen how this flexibility might be counterbalanced by the aid's capability to monitor the user. In a complementary fashion, the user might quite gladly hand over certain tasks to the aid, retaining, however, the capability of monitoring its performance and interjecting his own judgments where he deems it appropriate.

In a second project for ONE, DSC has developed decision aids which can incorporate both objective data and subjective judgment. [23][24][25] A special focus of this work has been the analysis of passive sonar data to estimate the range of a target on a nuclear attack submarine. This task, logically, should be included within the situation assessment feature of the attack planning aid described in previous sections. In particular, work on this aid has shed some light on how the Adjust Module might be utilized to facilitate user inputs into an otherwise automatic process.

Problem Setting

Numerous techniques are available for estimating target range - based on different aspects of the data (e.g., bearing, intensity, angle between direct and reflected sound paths) and using different analytical tools and assumptions. Typically, since their sources of error are both pronounced and different, they produce quite diverse estimates. Confronted with a divergent set of estimates, the commander is likely either to suspend judgment about range altogether or to focus on only one or two favored techniques, at the expense of others that might either corroborate or contradict them. Attack may be needlessly delayed while a good solution is improved, or be launched prematurely based on overconfidence in a bad one.

Pooling aid

A range pooling decision aid has been developed, utilizing a Bayesian framework to assist the command

staff in balancing and integrating the diverse sets of relevant information. The aid displays evidence (i.e., particular ranging techniques with assessments of their quality) as well as conclusions (a single best guess as to target range together with an interval of uncertainty). This aid has been implemented for testing purposes at the Naval Underwater Systems Center (NUSC-Newport).

For present purposes, two critical features of the aid should be noted:

- It can operate in a completely automatic mode.

Default estimates of pooling parameters, i.e., weights describing the precision of the solutions and their correlations, are based on at-sea exercise data. Ultimately, default parameters will be contingent on a variety of environmental and threat characteristics.

- The user can interpose his own assessments in addition to or in place of default estimates at any point in the range pooling process.

Preliminary Testing of Interactive Modes

The pooling aid has been tested in three modes:

- (1) totally automatic (default weights),
- (2) totally subjective (weights supplied by user), and
- (3) user override (default pooled solutions adjusted by user).

Prerecorded data from at-sea exercises were used to simulate conditions (2) and (3). Recorded command staff estimates ("system solutions") were used to derive subjective weights by multiple regression of command staff estimates on the particular ranging techniques. Command staff adjustment of default pooled solutions was simulated by pooling command staff estimates and default pooled estimates.

Figure 3 summarizes the results of this test for two different samples of Rangex data:




Figure 3. Ratio of Mean Absolute Error (MAE) for Various Interaction Modes to MAE for Default Pooled Solution. # = Rangex 1-78 data; * = Rangex 1-78 and 1-79 data

Subjective Pooling. Pooling with subjective weights was superior in accuracy both to the command staff estimate and to the specific ranging techniques. Although command staff estimates were superior to particular ranging techniques, the superiority of pooling with command staff weights to the command staff itself suggests that the information actually available to the command staff was not being optimally utilized by them.[26] These results would occur, for example, if the command staff were probabilistically selecting among estimates, with probabilities dependent on their relative accuracy, rather than pooling.

Automatic Pooling. Pooling with default weights was more accurate than pooling with subjective weights. This is not surprising since default weights were optimized for the type of data involved in the test. It is at least possible that subjective weights would outperform default weights in situations which differ sharply from exercise conditions.

User Override. The most accurate result was obtained in the third condition, where objective (default) data and subjective inputs were combined. This strongly suggests that, despite ineffective utilization, command personnel have access to relevant information not incorporated in the pooling aid.

This information can be tapped without burdening personnel with the task of formally pooling estimates. Leaving that job to the decision aid, appropriate staff might nonetheless monitor its performance and make adjustments when they observe significant discrepancies from their own intuitive solutions. The Alerting Module can assist in this monitoring, by alerting staff when default pooled range estimates or intervals of uncertainty fall outside a user-specified "plausible" region. Quite apart from any enhancement of user acceptance, our data suggest that incorporation of judgments in addition to objective data can improve the quantitative accuracy of aid outputs.

Other Applications

The requirement of stealth in warfare often imposes a severe constraint on communication among friendly units. Coordination can be achieved by prespecifying courses of action, but at the expense of flexibility. The combination of automatic aid functioning and user override capability offers a different approach. It may be applied to option generation (for example, user override of default option generation settings) and choice (for example, user override of default evaluations of outcomes, such as the relative worth of different types of targets). In either of these cases, default settings might be based on doctrine or mission directives; the provision of override will then set a balance between central guidance and flexible response to unique circumstances.

Conclusion

Both the attack planning aid and the range pooling aid have met with some success in initial demonstrations with representative potential users. Nonetheless, many if not most of the basic ideas presented in this report remain untested. Careful work remains to be done in delimiting the cognitive structures and modes of processing to which aids should cater, in defining and testing non-obtrusive prescriptive prompts, in identifying non-burdensome methods for incorporating judgments, and in developing guidelines to determine when and for whom methods of the sort described here are appropriate. The hoped for benefits include both increased user acceptance and improved system performance.

Note

* The work described was supported by the Office of Naval Research, Engineering Psychology Group, under Contract Numbers N00014-82-C-0138 (technical monitor J. O'Hare) and N00014-80-C-0046 (technical monitors M.A. Tolcott and G.S. Malecki), with the collaboration of the Naval Underwater Systems Center (Code 35). R.C. Bromage, R.V. Brown, J.O. Chinnis, L. Merchant-Geuder, J.W. Payne, and R. Pariseau have contributed in a variety of ways to this effort.

References

- [1] Mintzberg, H. Impediments to the use of management information. New York: National Association of Accountants, undated.
- [2] Sinaiko, H.W. Operational decision aids: A program of applied research for Naval command and control systems (Technical Report 5). Washington, D.C.: Smithsonian Institution, July 1977.
- [3] Cohen, M.S., Bromage, R.C., Chinnis, J.O., Jr., Payne, J.W., and Ulvila, J.W. A personalized and prescriptive attack planning decision aid (Technical Report 82-4). Falls Church, VA: Decision Science Consortium, Inc., July 1982.
- [4] Schum, D.A., and Martin, A.W. Assessing the probative value of evidence in various inference structures (Research Report 81-02). Houston, TX: Rice University, 1981.
- [5] Wason, P.C. On the failure to eliminate hypotheses in a conceptual task. Quarterly Journal of Experimental Psychology, 1960, 12, 129-140.
- [6] Wason, P.C. On the failure to eliminate hypotheses...A second look. In P.C. Wason and P.N. Johnson-Laird (Eds.), Thinking and reasoning, 1981, 7, 242-248.
- [7] Einhorn, H.J. Learning from experience and sub-optimal rules in decision making. In T.J. Wallsten (Ed.), Cognitive processes in choice and decision behavior. Hillsdale, NJ: Lawrence Erlbaum Assoc., Inc., 1980.
- [8] Lichtenstein, S., Fischhoff, B., and Phillips, L.D. Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press, 1982.
- [9] Schum, D.A., DuCharme, W., and DePitts, K. Research on human multistage probabilistic inference processes. Organizational Behavior and Human Performance, 1973, 10, 318-348.
- [10] Beyth-Marom, R. The subjective probability of conjunctions (Decision Research Report 81-12). Eugene, OR: Decision Research, December 1981.
- [11] Slovic, P., Fischhoff, B., and Lichtenstein, S. Accident probabilities and seat belt usage: A psychological perspective. Accident Analysis and Prevention, 1978, 10, 281-285.
- [12] Streufert, S., and Streufert, S.C. Stress and the measurement of task performance. I. Decision making in complex tasks (Technical Report 3). Hershey, PA: The Milton S. Hershey Medical Center, Department of Behavioral Science, 1981.
- [13] Driver, M.J., and Mock, T.J. Human information processing, decision theory style, and accounting

information systems. Accounting Review, 1976, 50, 490-508.

[14] Pitz, G.F., Sachs, N.J., and Heerboth, J. Procedures for eliciting choices in the analysis of individual decisions. Organizational Behavior and Human Performance, 1980, 26(3), 396-408.

[15] Payne, J.W. Information processing theory: Some concepts and methods applied to decision research. In T. Wallsten (Ed.), Cognitive processes in choice and decision behavior. Hillsdale, NJ: Lawrence Erlbaum, 1981.

[16] Johnson, E.J., and Payne, J.W. Accuracy and effort in choice. Pittsburgh, PA: Carnegie-Mellon University, March 17, 1983.

[17] Tversky, A. Elimination by aspects: A theory of choice. Psychological Review, 1972, 79(4), 281-299.

[18] Simon, H.A. Models of man, social and rational: Mathematical essays on rational and human behavior in a social setting. New York: Wiley, 1957.

[19] Svenson, O. Process descriptions of decision making. Organizational Behavior and Human Performance, 1979, 23, 86-112.

[20] Chase, W.G., and Simon, H.A. The mind's eye in chess. In W.G. Chase (Ed.), Visual information processing. New York: Academic Press, 1973.

[21] Klein, G.A. Automated aids for the proficient decision maker. IEEE Transactions on Systems, Man, and Cybernetics, 1980, 301-304.

[22] Cohen, M.S., and Freeling, A.N.S. The impact of information on decisions: Command and control system evaluation (Technical Report 81-1). Falls Church, VA: Decision Science Consortium, Inc., 1981.

[23] Cohen, M.S., and Brown, R.V. Decision support for attack submarine commanders (Technical Report 80-11). Falls Church, VA: Decision Science Consortium, Inc., October 1980.

[24] Cohen, M.S. Quantitative validity of objective and subjective inputs for a target range pooling decision aid. Proceedings of the 49th Military Operations Research Symposium (U) (MORS-49). Alexandria, VA: Military Operations Research Society, Inc., June 1982. (S)

[25] Bromage, R.C., Brown, R.V., Chinnis, J.O., Jr., Cohen, M.S., and Ulvila, J.W. Decision aids for submarine command and control. Phase III: Concept implementation (U) (Draft Technical Report 83-2). Falls Church, VA: Decision Science Consortium, Inc., March 1983. (C)

[26] Daves, R.M. The mind, the model, and the task. In Restle, F. et al., (Eds.), Cognitive theory. Hillsdale, NJ: Lawrence Erlbaum, Assoc., 1975, Vol. 1, 119-130.

AD POU 2893

COORDINATION OF BATTLE GROUP WARFARE COMMANDERS THROUGH SUMMARY DISPLAY TRANSFERS

Conrad J. Grant

The Johns Hopkins University Applied Physics Laboratory
Johns Hopkins Road
Laurel, Maryland 20707

Abstract

There are three primary methods for timely information transfer that are employed in the Fleet today. These consist of digital data transfer (such as Link 11), record message traffic (such as Fleet Satellite Broadcast or NAVMACS), and voice communications. With the advent of command facilities with advanced display and processing capabilities, there evolves a fourth method for information transfer: command summary display. The summary display can be composed of graphic information (such as Large Screen Displays) or alphanumeric information (such as the Automatic Status Boards). It has been determined that in many situations, the Composite Warfare Commander (CWC) for the Battle Group requires the same summary display information that a subordinate warfare commander has at his disposal in another command facility. Having similar information displayed at both sites enhances coordination between the commanders, and transferring this information in the form of summary displays can greatly reduce the need for replication of special purpose processing and data bases. While the transfer of summary displays can greatly enhance the coordination between commanders, it also places a heavy burden on link communications. Therefore, it is imperative that only essential information be transferred and that compression and abbreviated representation techniques be used in the transfer of the display information.

This paper will focus on the transfer of command summary displays between command facilities to enhance warfare commander coordination and the requirements that these display transfers place on the display, decision, and communications systems that are involved.

I. Introduction

Composite Warfare Commander Concept

Current Navy directives provide for the Commander Battle Force/Group at his discretion to employ the Composite Warfare Commander (CWC) doctrine in command of his forces.^[1] The use of this doctrine under control of the Officer in Tactical Command (OTC) establishes a CWC who is responsible for defending the force from air, surface, subsurface, and other types of threats. When necessary, the CWC will designate three subordinate warfare area commanders who, when delegated authority, coordinate offensive and defensive operations with conflicts being resolved by the CWC, who retains control by negation.

Command Support Facility Development

The Johns Hopkins University Applied Physics Laboratory (JHU/APL) under the Battle Group Anti-Air Warfare Coordination (BGAAWC) Program is developing the Battle Group Anti-Air Warfare Display Group (BADG) to determine the display, decision, and communication requirements of the Anti-Air Warfare Commander (AAWC). The BADG serves as a development tool for the design and

test of advanced display and decision aid capabilities for the AEGIS Display System (ADS) to support the AAWC aboard the new AEGIS Class cruiser (CG-47). The displays of BADG and ADS comprise four Large Screen Displays for graphic presentation of track data, ten CRT monitors configured as Automatic Status Boards for display of summary alphanumeric information, and multiple other special purpose displays. Once initiated, all displays can be automatically updated from the data base by the BADG processors.

In a parallel effort, the Navy Ocean Systems Center (NOSC) is developing the Tactical Flag Command Center (TFCC) and associated equipment to support the Battle Group OTC and CWC aboard Naval aircraft carriers and other selected combatants.

With continued development of these systems it has been recognized that there is an imperative need to determine interoperability requirements for the ADS and TFCC in order to fully meet their respective Commander's requirements for coordination and exchange of information. Effort was initiated at JHU/APL to propose a set of ADS/TFCC interface guidelines^[2] that discussed the factors that affect the interoperability of these two systems, the types of information that should be transferred to meet AAWC and OTC/CWC command and coordination requirements,^[3] and the communications, processing, and display systems that will be available to perform such an information transfer.

CWC Concept Premise

The development of interoperability requirements between afloat command facilities is dependent on the extent to which the CWC concept will be employed in the Battle Group (BG). In order to develop capabilities to meet the most stressing requirements for coordination, the following assumptions were made. First, the location of the warfare commanders on different platforms intensifies coordination requirements. Second, full delegation of responsibilities by the CWC to the warfare commanders is the most stressing case. Third, the maintenance of the warfare area complete track picture is the most communication intensive of the warfare commander's responsibilities.

II. Command Facility Capability

Assumed Capability

In general, the command facilities being developed to support the CWC and his subordinate warfare commanders provide similar capabilities in terms of communications, display, data base, and processing, although the methods of implementation may differ significantly. In developing the interoperability requirements, the following assumptions were made with respect to the command facilities capabilities.

Each command facility will have access to the full set of tactical digital data links and satellite links that are available to a surface combatant. Figure 1 shows the 1990 communications environment

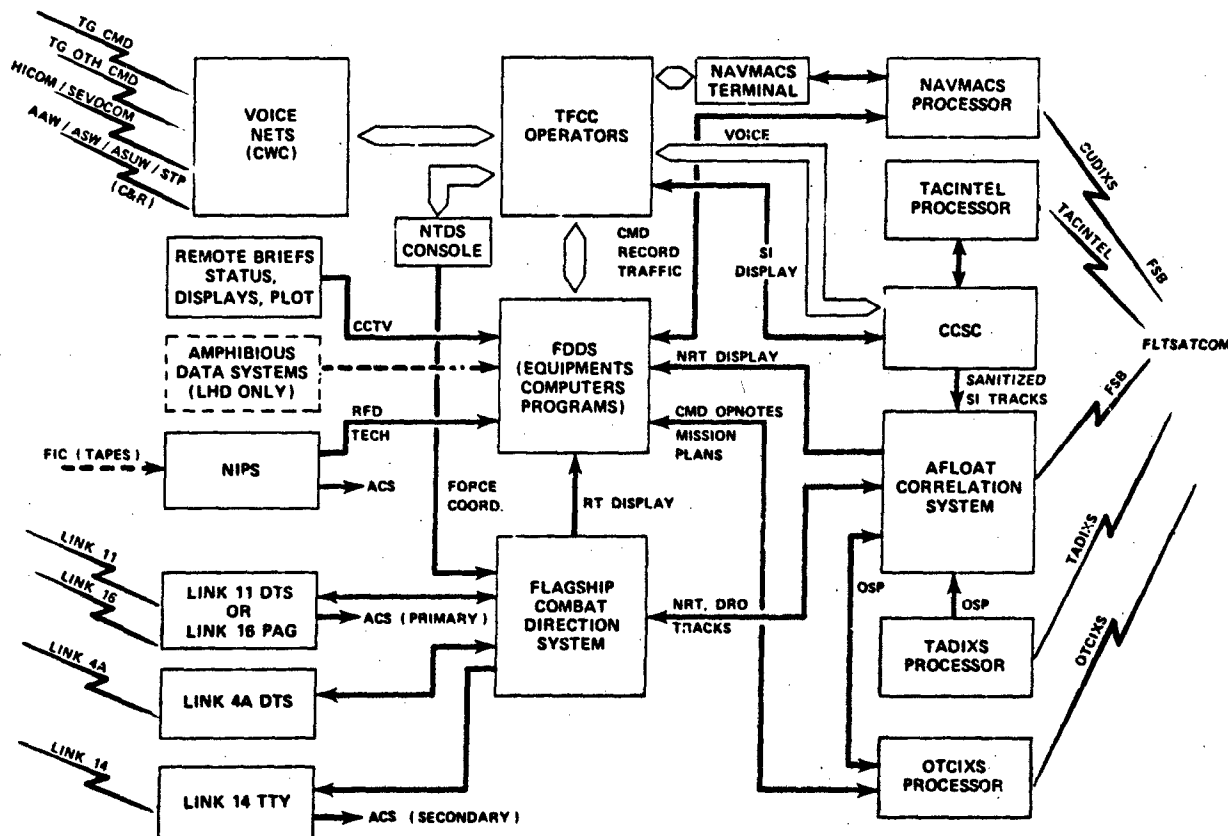


Figure 2. TFCC Communications Environment.

IV. Summary Display Transfers

Attributes

It has been determined that in many tactical situations, the CWC for the Battle Group requires the same summary display information that his subordinate warfare commanders have in their command facilities.^[1] For example, while under the CWC concept the warfare commander may be delegated the responsibility to maintain the force track picture for his warfare area as well as force status for sensors, weapons, and engagements, the CWC also requires this information to exercise control by negation.

Having similar summary information displayed at both sites enhances coordination between commanders when used in conjunction with other methods of communication (such as a voice discussion of a track picture). Transferring this information in the form of summary displays can greatly reduce the need for replication of special purpose processing and data bases. The summary displays can be sent automatically on alert, periodic, event, or query basis by the command facility processing system. This method of information transfer has a drawback though, in that it does not facilitate local storage (other than snapshot) or processing of the information by the receiving system. If the intent is to perform a data base to data base transfer, then a computer-to-computer data link transfer would be much more efficient.

While the transfer of summary displays greatly enhances coordination between commanders, it can also exact a great price in link communications loading. Therefore, only essential information should be transferred, and techniques for compression and abbreviation should be used to lessen the communications burden.

Automatic Status Board Transfers

An ASTAB contains up to 16 lines of 32 alphanumeric characters in multiple adaptable formats that, once initiated, are automatically updated and maintained by the processor. In order to transmit the ASTAB display, a Navy character-oriented message format called RAINFORM GOLD was chosen because of its capability to handle free text messages. The display contents are packed in the message in a full display compressed serial format. Figure 3 shows an example ASTAB and Figure 4 shows the corresponding RAINFORM GOLD message. In this example, an asterisk (*) followed by a number (n)

FORCE COMMUNICATIONS STATUS			
LINKS	STATUS	EMCON	JMD
JTIDS	GRN	U	GRN
LINK 11	GRN	U	YEL
LINK 4A	GRN	U	GRN
VOICE NETS			
A4W CMD	GRN	U	GRN
A4W RPT	GRN	U	GRN
LINK COORD	GRN	U	GRN
CAP CNTR			
G2R	GRN	U	YEL
F1B	GRN	U	GRN
R3Y	GRN	U	YEL
N7S	GRN	U	GRN

Figure 3. Example Automatic Status Board (ASTAB).

```

ZYN-UUUUU(I/O)
(PRECEDENCE)(DTG)(I/O)
(ORIGINATOR)(I/O)
(ADDRESSEE)(I/O)
(UNCLASSIFIED)(I/O)
MSGID/CMD/COMAND/GOLD/SER 0999/JUN(I/O)
NARR/ASTB:FORCOMSAT/_FORCE_COMMUNICATION_STATUS/(I/O)
/*6 LINKS*14 ST/TUS*4 EMCON*4JMD/(I/O)
/*7 JTIDS*14 GRN*7 U*6 GRN/(I/O)
/*7 LINK 11*12 GRN*7 U*6 YEL/(I/O)
/*7 LINK 4A*12 GRN*7 U*6 GRN/(I/O)
/*6 VOICE NETS/(I/O)
/*7 AAW CMD*12 GRN*7 U*6 GRN/(I/O)
/*7 AAW RPT*12 GRN*7 U*6 GRN/(I/O)
/*7 LINK COORD*9 GRN*7 U*6 GRN/(I/O)
/*6 CAP CNTR/(I/O)
/*7 GZR*16 GRN*7 U*6 YEL/(I/O)
/*7 FIN*16 GRN*7 U*6 GRN/(I/O)
/*7 R3Y*16 GRN*7 U*6 YEL/(I/O)
/*7 N75*16 GRN*7 U*6 GRN/END/(I/O)
ENDAT/DECL_31_DECEMBER_1986(I/O)
BT(I/O)

```

HEADER

Figure 4. Example ASTAB RAINFORM GOLD Message.

denotes a string of n spaces. This serial compression technique is economical for a string of 3 or more spaces between characters and can save up to 50 percent in message length for displays with few characters and many spaces.

A formatted character field transfer was considered, where each command facility would store a copy of all possible ASTAB formats with header and column information, and only tactically significant information fields would need to be transferred. This format was rejected since it constrained the type of ASTABs, required force ASTAB configuration control, and increased processing complexity.

Large Screen Display Transfers

A LSD contains graphic and alphanumeric information as well as multiple special symbols, that are commonly configured to represent force track plots. The RAINFORM GOLD message format was also chosen for the LSD display transfer. In this case the graphics information is represented in the message in the form of a display overlay. The overlay for a circular boundary, for example, is described by the letters "CIR" followed by the position of the circle center, and its radius. Figure 5 shows the overlay representation for a line, dashed line, circle, dashed circle, arc, rectangle, string of text, and a symbol in a RAINFORM GOLD message.

As previously mentioned, the information transferred from a warfare commander's LSD track plot includes only interpretative information and special tracks. It does not include maps and track position information which are assumed to already be available at both sites. Not only does this reduce communications loading, it also eliminates interoperability problems due to displaying map information through different types of map projections used in the command support systems (mercator versus orthographic projection).

```

HEADER
NARR/OVLY: TITLE/LSD DISPLAY NAME/(I/O)
/LINE/LAT LON/LA: LON/(I/O)
/DASH/LAT LON/LAT LON/(I/O)
/CIR/LAT LON/RADIUS(NM)/(I/O)
/CIRD/LAT LON/RADIUS(NM)/(I/O)
/ARC/LAT LON/RADIUS(NM)/ANGLE 1 (BAM)/ANGLE 2/(I/O)
/BOX/LAT LON/LAT LON/ROTATION(BAM)/(I/O)
/TEXT/LAT LON/CONTENTS/(I/O)
/SYMBOL/LAT LON/NAME/(I/O)
ENDAT/(I/O)
BT(I/O)

```

Figure 5. Large Screen Display, RAINFORM GOLD Formats.

Communications Media

It is assumed that the summary display RAINFORM GOLD messages will be sent within the BG via the Officer in Tactical Command Information Exchange System (OTCIXS) of FLTSATCOM or via a character-oriented subchannel of the Joint Tactical Information Distribution System (JTIDS). It is important that multiple paths exist for transmission of this information to improve survivability and throughput.

V. Conclusion

Communication between the Battle Group CWC and subordinate warfare commanders through summary display transfers can enhance coordination between commanders when used in conjunction with other communications methods. The display transfers require advanced processing, data base, and display capabilities, but can be performed using existing Navy message formats on existing and near-future communications links. The transfer of graphics displays requires the development of an intersystem graphics language such as the aforementioned graphic overlays.

Future efforts in this area include examination of other intersystem graphics languages such as North American Presentation-Level-Protocol Syntax (NAPLPS),^[4] modeling of command facility display message exchange to determine communications loading, and the actual demonstration of command summary display transfers between the BADG facility at JHU/APL and the TFCC facility at NOSC.

References

1. Naval Warfare Publication 8-1, "Composite Warfare Doctrine (U)", CONFIDENTIAL, dated June 1982.
2. C. J. Grant, "Interface Design Guidelines for BADG-TFCC Communications (U)", CONFIDENTIAL, JHU/APL Fleet Systems Report FS-82-114, dated November 1982.
3. F. C. Leiner, "Current Requirements for Information Transfer Among the AAWC, the OTC/CWC, and Other Command/Coordination Authorities within the Battle Group (U)", CONFIDENTIAL, JHU/APL Fleet Systems Report FS-82-275, dated December 1982.
4. J. Fleming and W. Frezza, "NAPLPS: A New Standard for Text and Graphics," *Byte*, Volume 8, Number 2, pp. 203-254, February 1983.

A MAN-MACHINE INTERFACE CONCEPT FOR A STATE-OF-ART, SHIPBOARD, COMMAND/CONTROL CONSOLE

Dr. Glenn A. Osga

Code 8231
Naval Ocean Systems Center
San Diego, California 92152

Existing tactical display/control consoles which are located in Combat Information Centers (CICs) aboard Navy vessels do not take full advantage of existing hardware/software/human factors technology. In addition to limitations imposed by hardware constraints, the user-interface software imposes a difficult-to-learn interface upon the complex task-demands of the command and control environment. The Naval Ocean Systems Center (NOSC) has developed a prototype for a command/control console which features many design advantages in comparison to existing consoles. The new configuration presents a diverse array of human-engineering issues, some specific to this console and others generic to all consoles. An overview of these issues and relevant research conducted at NOSC is presented in this paper.

Introduction

In 1978, NOSC began conceptual design of a CIC console which was originally conceived to meet the needs of light-surface combatants such as hydrofoils. The applications for this new console expanded to all Navy combatants as the design was determined to alleviate many present console problems. Specifically, these problems include: High power consumption (requiring water cooling), high cost, heavy weight, and lack of flexibility/growth. The design approach was to incorporate current computer and display technology to alleviate these problems. The console is currently called the Lightweight Modular Display System (LMDS), and it has shown to be very promising in terms of meeting these functional design goals.

LMDS Hardware Configuration

The present console configuration consists of two primary user-interface units, shown in Figure 1. A single, high-resolution, monochrome, 15-inch (diagonal) CRT display unit and an 11- x 15-inch high-resolution digitizer tablet are used. The display resolution is 1024 x 1368 discrete points or pixels. The effective viewing area is 8 x 10.75 inches. Figure 2 presents the general display format seen by the operator. The center 8- x 8-inch area is used primarily for tactical/positional and tabular data presentation. The display area on the right and left margins provide information such as control labels, control actuation feedback, system status, and other prompts. The area to the left of the display unit is reserved for communications equipment and analog display controls.

The digitizer tablet has discriminable points at .004-inch separation (250 per inch) in both X and Y dimensions. It responds to light pressure from either a finger touch or stylus, and transmits a digital word which represents the centroid of the pressed area. A cursor is presented on the display




Figure 1. Lightweight Modular Display Prototype Console

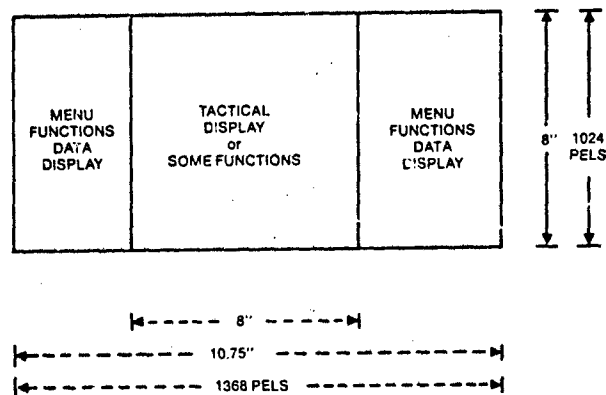


Figure 2. General LMDS Display Format

corresponding to the location on the tablet. Position is updated rapidly (approximately 100 msec.) such that position feedback delay is not perceivable. There are no labels printed on the tablet in the current configuration.

A Plexiglas cover is placed over the tablet to provide usable space for the placement of materials such as printed messages and notes. A movable fold-down keyboard is being considered which would store beneath the Plexiglas cover and fold down over the tablet. Plungers placed beneath the keys would activate the tablet. Metal overlays which would help guide finger placement on the tablet are also being considered.

Key Human Factors Study Areas

The interface design goals for LMDS include the following:

1. Improve feedback to user
2. Reduce user-memory load
3. Accommodate different user skill levels
4. Improve data filtering capability
5. Improve response time/accuracy for track designation
6. Reduce operator errors
7. Reduce operator training requirements

Meeting these design goals will require careful study of user information requirements, user tasks, user skill levels, operational environment and training methods.

User Requirements/Tasks. Since the LMDS console represents a radical departure from existing consoles it forces the designer to make a comprehensive study of the interface needs for all command/control tasks. User information needs and tasks will vary considerably across the personnel structure of the combat team. Whereas previous consoles were primarily used for 'input' operators (i.e., those tracking targets and updating tactical data bases), future consoles will also be used as workstations for high-level decision-makers. In addition, many of the tasks which are currently done manually will be automated to a greater extent in future Combat Direction Systems (CDS). Thus, more tasks will be delegated to a supervisory control mode with the operator requiring frequent updates on automated task processing with the capability for manual override. Information requirements for higher-level decision-makers will be increased by an enlarging of the threat area and sensor ranges, coupled with an increased variety of countermeasures/weapons. Users must be able to cope with an overwhelming amount of available data by using display controls that allow for data filtering of positional and tabular data.

User Skills/Training/Environment. It is useful to think of Navy console users as falling along a continuum from novice to expert users. Novices are new to the Navy Tactical Data System (NTDS) by virtue of being new recruits or by transfer from a non-NTDS ship. Typical training involves a 3- to 4-week 'input' course followed by a 3- to 4-week 'User' course for higher-level decision-makers. Training, therefore, can be described as a quick overview of the 1000-plus command vocabulary of the NTDS. The

typical console user graduates to on-the-job training where he acquires a working knowledge of the system. This cycle is constantly repeated as personnel are transferred or leave service. These conditions clearly dictate the need for an interface which can prompt and aid the novice user, while providing flexibility and 'shortcuts' for the experienced user. Present consoles and interface software do very little of either.

The shipboard environment requires that equipment meet tough militarization requirements, and that noise, vibration, and platform movement are considered in control/display design. Watch lengths are usually 4 hours, but may be longer, and the sea-state may impose a degraded environment for control and display. These constraints must be considered together with the precision required when selecting input devices (i.e., fixed or movable). Equipment and communication 'noise' combine to produce a high level of ambient sound. Ambient noise will limit the effective use of aural feedback.

Human Engineering the LMDS Configuration

A comparison of user tasks and information requirements with the hardware configuration provides a 'shopping list' of Human Factors considerations from which trade-off studies can be generated. Table 1 presents a list of the major task areas and a subjective rating concerning the adequacy of the display and tablet in supporting each task area. Performance within each of these task areas is under continuing study at the NOSC LMDS Laboratory.

Table 1: Overview of the LMDS Configuration Support for Command and Control Task Areas

User Tasks/ Job Requirements	LMDS CONFIGURATION	
	Digitizer Tablet	Single 15-in. Display
Track Designation		
Cursor Positioning	adequate	adequate
Alphanumeric Data Input	poor	adequate
Alphanumeric/Tactical Data Display	--	limited
Menu/Function Selection	adequate	adequate
Procedural/Interface Tasks	software dependent	software dependent

Track Designation/Cursor Positioning. The operator must be able to designate individual tracks in order to perform some function (identification, interrogation, data amplification, etc.). He must also be able to position a cursor to select functions, menu items, and place graphics at precise locations upon the display. These tasks are performed by using a trackball device on existing consoles. A recent NOSC study compared a digitizer tablet with a 1:30 control/display (C/D) ratio to a trackball device for a cursor positioning task.[1] This study indicates that the tablet device will yield comparable performance to existing consoles for cursor positioning tasks.[2]

Alphanumeric Data Input. Requirements for alpha character data entry vary across operator positions and ship types. For example, operators may add to the graphic track symbols. Certain higher-level AEGIS ship-class operators may enter

doctrine data through a 3 x 6 button array of alphabet characters. The operator must page back and forth between two arrays which each contain half of the alphabet. The paging has shown to be cumbersome, but practical, for minimal data entry. Indications are that future requirements for alpha data input will increase for tasks such as doctrine entry, mission planning parameter entry, and decision-aid use. Numeric data are entered through a numeric keypad called the Digital Data Entry Unit (DDEU). This device is used for a variety of tasks as tracks are commonly referred to by their four-digit identification number. The digitizer tablet will present special problems for alphanumeric input tasks due to the lack of tactile feedback and limited possibility for auditory feedback due to background noise. Tactile feedback may be provided in the form of 'locator bumps' on the tablet; however, this may aid function selection tasks much more than data input tasks. Errors of character omission and double entry will likely be much higher than a conventional keyboard. Studies are planned to investigate the need for a fold-down keyboard for alphanumeric data input.

Alphanumeric and Tactical Data Display. Modern C² systems can quickly overload operators with irrelevant and densely formatted data. A basic design question is how many displays are needed and how to appropriately format and combine alphanumeric and graphic data. Current state-of-art consoles usually include the Planned Position Indicator (PPI), or radar screen as the primary display supported by an adjacent CRT for alphanumeric data. This configuration is based more upon engineering feasibility than consideration for what the operator needs to perform his job. Digital radar and emerging scan conversion techniques make possible the conversion of alphanumeric and sensor (tactical) data on a single display. Techniques such as windowing and scrolling/paging may permit a single display to provide all the data needed by the operator. The feasibility of a single display will be considered in future NOSC studies. A recent study of tactical symbology fonts for raster displays established a baseline size for raster-drawn symbology.[3] A current NOSC study is evaluating different symbology font styles which are derived from this baseline font. The objective of these studies is to produce symbology which is 'compact' and minimizes display clutter while being readable at adequate error levels.

Menu Formats/Function Selection. A diverse array of design questions are included in this topic area. First, the C/D ratio must be optimized for fine and gross cursor positioning tasks. The optimum ratio will determine minimum tablet size. The menu formats and method of function selection are related issues which will require careful study. Function selection methods include: single action, single action with software delay, double action, multiple actions. Major considerations in method selection are the severity of an input error, desired response speed, isolation of the function on the display page, and multiple sequence tasks. The guiding rule for method selection is minimize the number of alternate methods and be consistent in their application.

General Interface Issues. A number of diverse issues fall under the general category of tasks related to the overall operating structure. Feedback requirements include: input verification, process completion verification, state-of-system, and error messages.[4] User-feedback has been most noticeably neglected in present-day Naval combat

systems. Given the training structure defined earlier, the system needs to accommodate novice and expert users. Tutorials, training modes, and menu by-pass features are needed for various user skill levels. Programmable default options with edit capability should be available to users. Users should be able to 'verify to proceed' rather than be required to fill in menu selections. Multiple default lists would be extremely valuable for tactical symbology and display graphics. The menu network structure must be compatible with user expectancy and should minimize display paging and search time within a given frame.

Conclusions

This paper has touched upon a number of human factors issues which arise in the design of a multi-purpose command and control console for emerging combat direction systems. NOSC will be conducting an ongoing research program to address a variety of interface issues, some specific to the current LMDS project, others applicable to a generic console. Operator needs are often neglected in current command and control interfaces, most notably the lack of proper feedback, and complex procedures which require extensive learning while placing a heavy demand upon operator memory. With a comprehensive human factors program, many interface problems can be addressed early in the console/combat system development cycle, to ensure that past design errors are not needlessly repeated.

References

1. Davenport, E.W. and Bemis, S.V. Tracking Performance; Trackball vs. Digitizer Tablet. Technical Report 824, Naval Ocean Systems Center, San Diego CA, 1982.
2. Albert, A.E. "The Effect of Graphic Input Devices on Performance in a Cursor Positioning Task." Proceedings of the Human Factors Society 26th Annual Meeting, 1982, 54-58.
3. Onga, G.A. An Evaluation of Identification Performance for Raster Scan Generated NTDS Symbology. Systems Exploration, Inc., Report, Contract No. N00123-80-D-0263, San Diego CA, June 1982.
4. Gould, J.D. "Man-Computer Interfaces for Informational Systems: Lecture to Human Engineering Short Course," in Huchingson, R.D., New Horizons for Human Factors in Design, McGraw-Hill, 1981.

SECTION III

Surveillance: Issues and Algorithms

Preceding page blank

MULTI-SENSOR FUSION, COMMUNICATIONS AND INFORMATION WARFARE

Daniel Schutzer

INTRODUCTION

Today's fusion problems are chiefly concerned with organizational and procedural issues. The technology they employ is mostly available state-of-the-art. The future brings a new set of concerns centered about issues that are more technical in nature. Future military command and control and weapons systems will likely be more distributed, more automated and smarter. They will probably include an advanced form of information warfare where sensing, information exchange, jamming, deception, and misinformation will be capable of being managed and orchestrated from a total mission objective perspective. As a result, the future fusion process will be required to handle and process on orders of magnitude increase in the volume and diversity of input data, faster. It will need to produce a great variety of information to feed automated C2 and weapons systems data bases through more interactive and responsive interfaces than exist today. At the same time it needs to analyze this data at a deeper level of understanding than ever before, scrutinizing and drawing inferences and conclusions about ones adversaries underlying beliefs, readiness, intentions and future actions from what is often times a suspect and spotty data base. Finally, these conclusions and inferences need to be presented in a clear, concise, honest, but convincing and timely manner. This paper presents a unified framework from which the necessary information may be fused, managed and presented to support command in such a future information warfare environment and discusses the associated technical challenges. This paper reviews various ongoing research programs that are addressing these challenges.

INFORMATION WARFARE CONTROL STRATEGY

When viewed in this manner, the strategy illustrated in Figure 1 emerges as a reasonable information handling systems design and control scheme. The objective of this control strategy is to maximize a measure of the combined state of knowledge of own forces and state of ignorance of opponent forces such as $(-u)v$, through judicious application of the information manipulative functions. A mathematical expression that represents these relationships has been developed in references (2), (3), and (4). It takes the form of the following differential equations:

$$\begin{aligned} \dot{u} &= -au - b(-u)v + c'(-u)(-v) + d(-u) \\ \dot{v} &= -a'v - b'(-v)u + c'(-v)(-u) + d'(-v) \end{aligned}$$

where

u = entropy, or ignorance, of friendly
 v = entropy, or ignorance, of enemy
 \dot{u} = change in entropy of friendly
 \dot{v} = change in entropy of enemy

$a, b, c, d, a', b', c', d'$ are positive constants.

Friendly sensor measurements, intercepts of enemy communications and sensor radiations, and human intelligence reports are received, combined, and interpreted to form two pictures; friendlies state of knowledge of the situation and the opponents state of knowledge. This includes such information as what units are located where, their course, likely destination and intentions, with indications of the confidence associated with these estimates. Since information warfare is a two-sided contest, it is assumed that both parties will seek to select that combination of information manipulative functions which serves their best interests. Accordingly, friendly should select that set of information manipulative functions which results in a maximum objective function assuming that the opponent will always select a corresponding set of information manipulative functions designed to minimize the same objective function. This is a worst case approach. It might prove desirable to provide the decision maker with other possible control strategies and their possible outcomes; such as one based upon a best case analysis and another based upon a most likely (derived from historical precedence) enemy response.

Of course the end objective is to prevail in battle, maximizing the enemies losses and minimizing ones own losses. It is shown in reference (2) that the amount of knowledge, or information, a combatant possesses directly influences his combat effectiveness. Accordingly, it is recommended that prior to the selection for execution of any set of information manipulative functions, a projection of battle outcome be determined which includes the likelihood of enemy initiation of conflict and the predicted losses and associated risks. This battle outcome prediction should be presented along with information manipulation control strategies so that a course of action may be determined which includes consideration of the initiation of conflict as well as pure information warfare actions.

FUSION PROCESS

We have discussed how, in the future there will be more sophisticated and capable systems providing a greater diversity and quantity of more timely data for fusion and analysis. The enemy, however, will be more sophisticated and capable in deception, camouflage and decoy. Consequently, we will be

Preceding page blank

required to make use of this increased data just to accurately monitor the enemy without being surprised or fooled. Moreover, the increased tempo of battle and the needs of the unified information warfare strategy, require us to analyze, process, and fuse, this increased volume of data in even shorter time. All this highlights the need for an improved more responsive fusion process.

More specifically, we are concerned with the following sorts of improvements:

1. The percentage of data received that is profitably used (exploitation of current data).
2. The volume of data that can be processed in a fixed amount of time (system capacity).
3. The ability to select and prioritize input data to be processed and/or discarded by its value added (selective processing).
4. The speed at which data can be received and processed (timeliness & throughput).
5. The quality of the information output (information credibility, reduction of ambiguity, resistance to camouflage, cover and deception, confidence in predictions).
6. The effective use of the historical archival data (exploitation of archival, historical data).
7. The behavior of the system under data overload/saturation (system stability).
8. The manpower requirements (number, skill level, training, fatigue).
9. The relevance of the processed information to the situation and the decision-makers needs (information tailoring, responsiveness/relevance, anticipation of information needs, and data reduction).
10. Effectiveness of the sensor collection tasking.
11. The integration of analysis products produced at different contributing sites/organizations and locations and the resolution of collection tasking conflicts and priorities amongst these various locations (analyst inter-site integration).

Our primary interests are with technologies concerned with the processing, organization, retrieval, presentation and dissemination of data and information. Deeper interests lie in gaining a greater understanding of the theories of human cognition and reasoning and of plausible inference and deduction as applied to the intelligence analysis process. At all levels of this process there is a need for better ways to process, reduce, and interpret data in order to produce timely and relevant information.

The technical challenges vary with the stage of the intelligence analysis process and the form of the data at that stage. It includes such areas as high speed signal processing and pattern recognition, symbolic representation and computation, extremely high density storage, image, speech and text understanding, distributed hypothesis generation and problem-solving, planning and scheduling.

The stages of the intelligence process are of course highly interactive in that data does not flow neatly left to right as implied in Figure 2, or in any fixed order. Rather, an intelligence analysis problem gets worked simultaneously at all stages or levels (top-down, bottom-up, and from the middle) with tightly coupled feedback loops between all stages or levels. Further, analysis is inherently a human-directed process where the information processing techniques developed would serve as an assistant, amplifier, and advisor to the intelligence analyst, and not as a replacement. Accordingly they should be designed to allow an analyst to do his job better, consider more variables, and work faster and more efficiently. Because of these attributes an architecture is proposed that borrows from the Hearsay III Blackboard model (a federation of experts who collaborate on solving a problem through the mutual sharing and updating of a common hierarchically structured multiple view of the evolving problem state). Key components of this architecture is illustrated in Figure 3 and is explained below.

FEATURE/PARAMETER EXTRACTOR

Sensor data comes initially processed and in multiple forms: formatted messages with features and parameters identified, narrative text, and imagery. In some cases even the pre-processed signal video (telemetry, ELINT waveforms and acoustic lofogram data) is available. It is assumed that all ELINT data gets processed and reduced to a message (fixed format or narrative) at the collection/sensor site. The raw imagery and signal video, provided would be in addition to the preliminary processing at the collection site and to be used for more in-depth analysis. Thus, when data enters it goes two places: initial feature parameter analysis; and data storage and retrieval. The feature parameter analysis function is primarily a pattern recognition syntactical and statistical analysis process where features and parameters such as PRF, modulation parameters, harmonic pairs, image ribbons and texture measures get reduced to various combinations of discriminant categories and classes which are then sequentially tested against an existing parameter/feature data base constrained by associated collateral data such as candidate object locations and characteristics, and time of observations. The objective of this sequential test process is to reduce, through elimination, contradiction and deduction possible associations between received feature and parameters to as small a set of feasible candidate objects, events and event hierarchies as possible. The result of this process produces mixed results. Some discriminants have very strong unique associations with a specific object or event, others can only be ambiguously or probabilistically associated with one or more objects and events. There are several technologies that would be useful in this area: high speed signal processors (e.g. VHSIC, array processors, systolic processors) would be needed to execute the compute-intensive multidimensional pattern and cluster analysis algorithms required against the potentially high volume of input data. Natural language processing techniques are needed to reduce the time and manpower currently needed to interpret the large volumes of narrative messages into formats suitable for further computer processing. The raw sensor data is assumed to be first stored. It is retrieved and processed only in response to special directed analyses and questions that get identified further down in the analysis process. Discussions as to how this data gets retrieved and processed will be discussed in

greater detail in the section on Data Storage, Retrieval and Inference.

INTERPRETER/CONCEPT GENERATOR

The Interpreter/Concept Generator attempts to logically relate the data pertinent to the current situation/state description to logical chunks of information related to, or associated with, concepts, objects events and event hierarchies. This logical association can probably be most naturally expressed in the form of frame/script based hierarchies. Each object will have slots representing attribute data such as physical features, electronic order of battle, current location, past locations, speed, and associated events, event hierarchies and rules of behavior. Events will include such items as "missile firing" or "report back". Each event chunk (or concept) will have slots associated with such attributes as measured observables, action taken, involved objects, location and time. As mentioned earlier, events can in turn be associated with other events as an identifiable time and/or order dependent sequence hierarchy; e.g. "coordinated attack operation". As data is placed in the data base the Interpreter attempts to fit the data to slots in the various object and event frames/scripts. Each object and event frame have Lemons (embedded procedures) that get triggered as a function of the specific data clues or operations to look for inconsistencies, establish linkages or alerts, perform special computations and take other actions. The Interpreter/Concept Generator maintains the common hierarchically structured multiple view of our collective knowledge of the evolving problem state. It should be capable of linking together one or more parameters, objects, and events with an associated confidence and propagating that confidence measure up the common multiple view hierarchy and across a chain of logical deductions.

HYPOTHESIS GENERATOR

The hypothesis generator is concerned with higher level aspects of the current situation under study such as enemy strategy and intent. It attempts to associate or infer what is the purpose or motivation behind the observed sequence of events; e.g. exercise, heightened state of readiness, a planned deception some other operation. This is one of the more difficult areas to deal with and is one of the more highly intuitive and unstructured functions. Accordingly it is envisioned as being the most manually-driven of the functions discussed so far. In the Feature/Parameter Extraction function and in the Interpreter/Concept Generator function, the data, knowledge representation and inferencing structure will likely be defined in advance with the process being fairly automatic; the analyst interacting on an exception basis.

In contrast to this approach, the formation and the selection of hypotheses is expected to be a dynamically changing function assisted by machine prompts, suggestions and aids, but closely controlled and managed by an analyst. For example, the Hypothesis Generator should be able to generate canned predictions of event sequences and object inter-relationships for several pre-specified enemy plans and strategies as a function of the environmental and situation-specifics. These initial event sequence predictions are anticipated as merely being a departure point for the analyst who should

have the capability, at a fairly high level of abstraction, to modify and adjust the initial hypothesis to a more suitable one of his choosing. The hypothesis generator should be capable of reevaluating the feasibility of two or more hypotheses simultaneously. At least one hypothesis should act as the duty skeptic, constantly checking for the possibility of deception, looking for the notable absence as well as the receipt of confirming data. This is something the human finds very difficult to do. Psychological studies reveal that a human typically forms a single hypothesis, and once formed, he tends to look for confirming evidence only. The incoming data and previously derived evidence should be continuously reviewed for indications of some change to any of the set of all possible hypotheses and it should alert the analyst when any significant change in the rankings of any of the hypotheses occurs; a previously likely hypothesis becomes suspect; a new hypothesis other than the ones currently under serious consideration is suggested; or perhaps, when there is not good fit of any hypothesis to the evidence at hand. In addition, the hypothesis generator should have both pre-set heuristics and user - specificable and adjustable heuristics that examine the incoming evidence for clues of confirmation or denial of currently favored hypotheses, and to assist in the nomination and formulation of new hypotheses. These heuristics should make use of such indicators as activity patterns and statistics, the observed time, order, and sequence of events, user-specified interesting and unusual associations or unexpected trends, negative information, and recognition of inconsistent patterns and inferences. These heuristics can generally be expressed as an unordered collection of "IF, THEN" conditional statements or production rules. Thus, the hypothesis generator is envisioned as borrowing heavily from the expert system production rule methodology. It should be capable of operating both in an automatic data-driven mode, alerting the analyst when some change of interest has occurred, as well as being able to be interrupted, tasked and/or queried by an analyst in more goal-directed search mode. These predicted activities are transmitted over a cooperative analyst network for confirming evidence in support of the generated hypotheses. In addition, the hypothesis generator exchanges alerts over the analyst net whenever unusual or highly interesting alerts, situations, or indications arise. The hypothesis generator is used to drive the collection manager/test planner to prioritize the order in which incoming data is processed directs the browsing through of historical data.

COLLECTION MANAGER/TEST PLANNER

The collection manager/test planner is concerned with assisting the analyst in providing appropriate feedback and tasking to the collection sensors and sources. The analyst has the opportunity to resolve inconsistencies, remove ambiguities, and to decide between hypotheses by appropriately cuing, alerting, and tasking the collector for confirming or predicted data. To provide this capability the collection manager/test planner maintains a data base concerning the various collection assets, their availability and appropriate physical characteristics. Both the interpreter/concept generator and the hypothesis generator signal the collection manager/test planner for assistance whenever they encounter a conflict or ambiguity that requires resolution or whenever they have reached a dead-end in their analysis. The

collection manager/test planner produces, based upon the set of initial hypotheses, a decision tree of collection measurements from which a best set of measurements is chosen. The best set of measurements is that set which will bring the system closest to reducing the candidate hypothesis to a single one under such user-specified constraints as the cost and likelihood of the collection tasking request being satisfied and its probable outcome. These conditional probability outcomes are determined from both rules previously provided by the analyst or approximately by the output of the predictor/simulator, discussed next. The collection manager test planner can be asked to provide a second or third best collection strategy, or can be interrupted and asked to provide a best collection strategy based upon a new set of constraints supplied by the analyst. Once a collection strategy is selected by the analyst, the appropriate tasking messages will be automatically generated and submitted for analyst approval prior to their submission to the appropriate collection tasking organizations. These same collection tasking requests will also be transmitted across the inter-analyst net to other cognizant analyst organizations working in related areas for coordination and for internal resolution of tasking conflicts and priorities.

PREDICTOR/SIMULATOR

The predictor/simulator is a fast simulation model that simulates/models the motions, behaviors, emissions and detections of target objects and sensors in a faster than real time mode. It is input with high level, fairly abstract descriptions that are outputted by the interpreter/concept generator, the hypothesis generator and the analyst. These inputs should be able to be viewed and modified by the analyst. The output of the predictor/simulator would be to move forward in time and to predict the measurements that would result from candidate sensor/collector taskings and their associated probabilities. Applicable technologies are the use of multi-processors to achieve the desired simulation execution speeds and of rule-based descriptors of the objects and events to allow for sufficient simulation flexibility and for good user interface. The RAND Rule Oriented Simulation System (ROSS) is a candidate technology for system. It has provisions to include rules of logic and behavior of the objects, models of the units, and sensors capabilities and movements, as well as models of the environment.

MATCHER/COMPARER

The matcher/comparer compares the incoming signal data at both the feature/parameter and the object/concept level with the outputs of the predictor/simulator. Either it finds matches, or it attempts to adjust, within the allowable bounds of the candidate hypotheses, the inputs of the predictor/simulator with respect to object orientation, motion and behavior in order to achieve a match. If a match is achieved, the appropriate portions of the common multiple view representation of the situation under study is updated appropriately. If no match is found, this fact and the amplifying data is reported to the analyst for further action. Much of this system is anticipated as being highly quantitative and computation intense VHSIC technology and advanced pattern matching algorithms appear appropriate to meeting this system requirement.

DATA STORAGE, RETRIEVAL & INFERENCER

The data base is envisioned as being extremely large, 10^{14} bytes or greater. It will contain a mixture of data and information types and forms including: historical narrative data, imagery and formatted computer data including hierarchical multiple views of the past, current and predicted state of objects, events and missions under analysis. The size and complexity of such a data base dictates a hierarchy of advanced storage media, such as magnetic tape, optical disc, and random access memory. Associate processor technology, array processors, systolic arrays, or the like may be required for rapid content-addressable data retrieval and update ("smart" memories). Heavy use is anticipated of the technologies of predicate calculus and deductive inferencing so that information may be queried and retrieved that is not explicitly stated but that may be derived or deduced from data that is explicitly stated and stored. Aids should be provided to allow the analyst to selectively browse through a large text or imagery data base. The analyst should be able to identify selected subjects, keywords, sounds, spatial features and relationships of interest, and the context under which they are of interest. The data storage, retrieval and inferencer should then be able to select relevant text, images, etc. for presentation to the analyst with a high confidence that all relevant items and only relevant items will be retrieved. If the number of items selected for retrieval exceeds some reasonable threshold, the analyst will be informed of the number of qualifying items. He may then request to see them all one-by-one, or he may choose to add additional qualifiers and constraints to his request for data. To satisfy this requirement, the system needs two capabilities: namely, the ability to use the situation peculiar context to translate the fairly high level conceptual information requests to an equivalent set of data indices and pattern feature matches and then to scan through large volumes of historical textual and imagery data for matches of the appropriate combinations of keyword, spatial features, and patterns at a speed commensurate with the transfer rates of the storage medium. This requires some impressive lower level feature and pattern extraction processors interfaced to a fast symbolic processor. (Reference (5))

DATA CHECKER AUDITOR

Closely related to the data storage, retrieval and inference is the data checker and auditor. This system acts as an independent data auditor constantly checking the collective data for internal consistency and plausibility. It should look for contradictions, exclusions, and negative evidence. This system is also concerned with data compaction and housekeeping; e.g. the merging, purging, and forgetting of data/information. To achieve this sort of capability it is anticipated that this system design will borrow heavily from the ideas of Truth Maintenance and theorem-proving and plausible reasoning.

PRESENTER

The presenter is envisioned as providing the man-machine interface, both locally to assist the analyst, and externally to provide the user his information needs and to support communications between the analyst community. As noted earlier, this inter-analyst network is needed to allow a more

effective integration and coordination of the products of the analysis community; e.g. means to better distribute the intelligence analysis load and to share the intelligence product, as well as to better coordinate and resolve conflicts and priorities in collection taskings related to this common distributed problem-solving activity. The presenter should make use of such advanced features as natural language and speech input and output, menus, graphics and explanation facilities to explain the rationale and evidence and line of reasoning behind any intelligence product. But it should do more than this, it should maintain a model of the user and how his information needs vary, dependent upon the situation. It should use this model to alert the user to situations of interest as they occur, and to reduce and to tailor the needs in support of his current decisions as opposed to drowning him in an information overload situation. The system should anticipate his need for additional further information. This system is what makes the fusion center appear responsive, timely, and relevant to the user without overloading or overburdening him with more data, at a rate faster than he is capable or cares to absorb (human bandwidth is approximately 10-40 Hz). There are techniques whereby, once this capability is in place that it can be used to reduce the communications overload normally associated with data base query and update (reference (6)). Finally, multi-level security will have to be addressed. We need to develop processors fast enough to handle the multi-level security overhead without significant performance degradation resulting.

REFERENCES

1. "An Analysis of Communications in Support of Command and Control", D. Schutzer, National Telecommunications Conference, November 27-29, 1979
2. A Dynamic Model for C3 Information Incorporating the Effects of Counter - C3 P. Moose, Dec 1980, NPS-62-81-025 PR
3. Concepts and thoughts concerning a control strategy for conducting information warfare, D. Schutzer 5th MIT/ONR Workshop on C3 systems, Aug. 1982
4. MORS Monograph on Command and Control, editors Schutzer, Shere, Huang, Vena, 1982
5. Concepts concerning the Management of Very Large Diverse Files, D. Schutzer EASCON 82, Sept. 1982
6. Applications of Artificial Intelligence to Military Communications, D. Schutzer, MILCOM 83, Oct. 31, 1983

FIGURE 1: INFORMATION WARFARE CONTROL STRATEGY

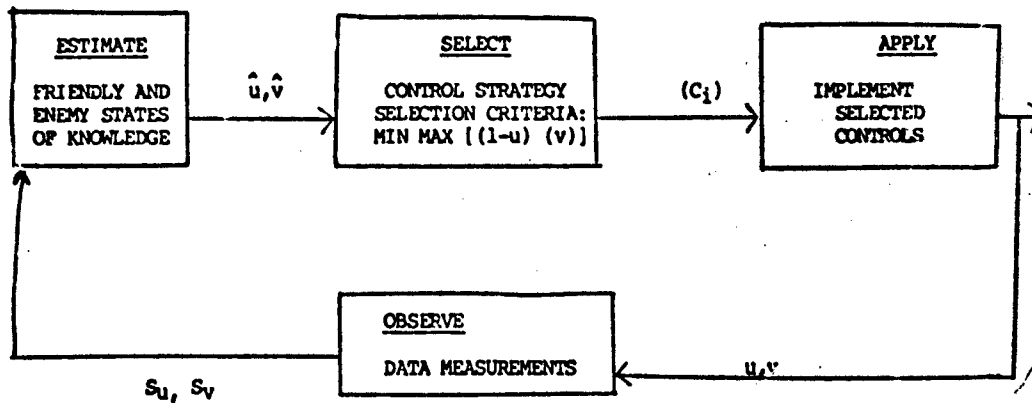


FIGURE 2 FUSION PROCESS

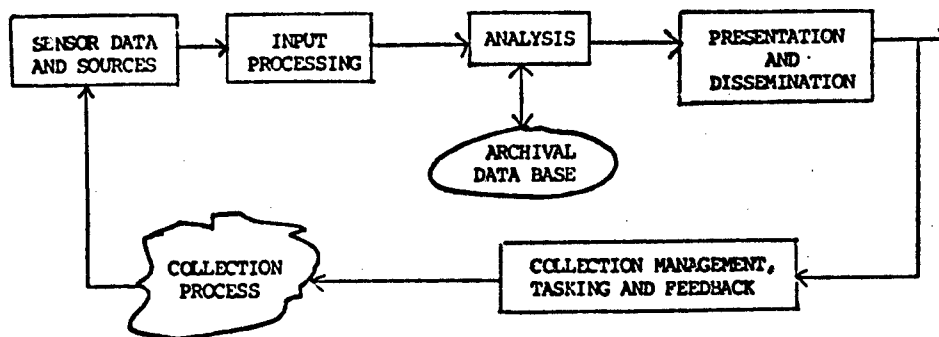
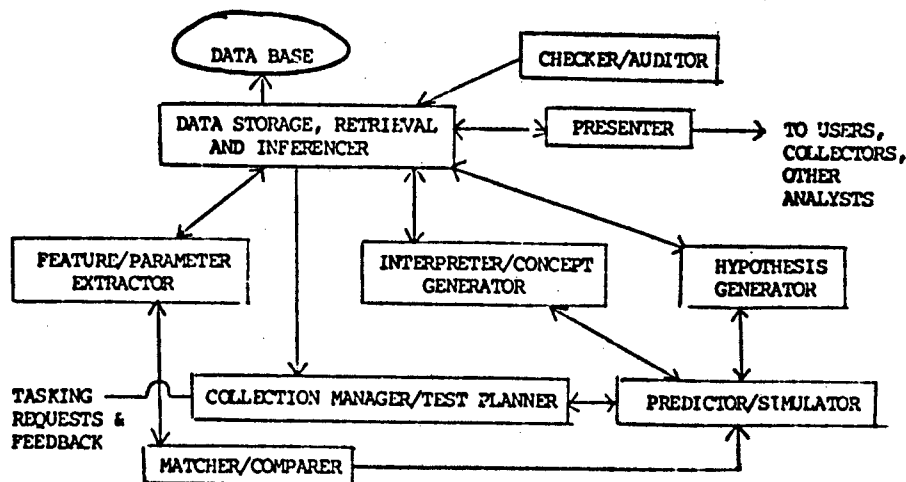


FIGURE 3: PROPOSED ARCHITECTURE



TECHNIQUES FOR DETECTING COVER AND DECEPTION

Cover and
Deception (C & D)

Robert P. Goldsmith and Ralph F. Gerenz

Betac Corporation
76 Treble Cove Road
Billerica, Massachusetts 01862Summary

The increasing sophistication of intelligence collection and analysis systems has given US decision makers a powerful tool to evaluate the actions and intentions of our potential adversaries. At the same time, however, these advances have in some respects increased our susceptibility to the skillful use of cover and deception techniques. Throughout history, the potential success of C & D operations has been determined solely by the skill of the practitioner, regardless of the sophistication of the intended victim in conducting C & D operations. Today, we face in the Soviet Union a nation which has both recognized the importance of C & D and has over the years demonstrated an impressive capability to deceive and mislead both its intended victims and the US and its allies. This paper outlines the salient characteristics of C & D, Soviet doctrine and application and some of the techniques which could be used to uncover cover and defeat deception.

to manage perceptions during peacetime, crisis and war. Similarly, C & D is a major concern in successful application of arms control and to the maintaining the validity of a deterrence policy. The operational context for C & D is shown in Figure 1.

The opposing decision-maker selects both an operations plan and a C & D plan, the first to achieve his objectives, and the second to manage friendly perceptions in such a way that counter-action will be misdirected or mistimed. These plans are executed as a course of action (1), aspects of which can be observed if friendly collection assets are present and active when they occur (2). Threat assessment (3) attempts to correlate information collected with a concept of how opposing forces would be used under a variety of circumstances, to produce an evaluation of what the opposing forces are trying to do. The friendly decision maker then selects a response (4), based on set policy and the capabilities of his own assets.

The Operational Context for Cover and DeceptionWhat is C & D?

"Cover" denies an adversary the intelligence data needed to plan and carry out operations, and it includes both camouflage and avoidance. Camouflage can be either passive, in which case it attempts to make the threatening activity appear either benign or not appear at all, or active, in which threatening activity is simulated where it does not, in fact, exist. Avoidance exploits knowledge of the adversary's collection capabilities and operational use to deny reconnaissance opportunity.

"Deception" seeks to use both camouflage and avoidance, together with genuine but misleading activity, to manage an adversary's perception of events, capabilities and planned actions.

A skillful user of C & D seeks to provide an adversary with pieces of information which appear genuine in themselves, and which fit a course of action which the adversary would find reasonable. In this, the C & D practitioner attempts to exploit the anchoring bias of the cognitive process [1], by presenting the strongest indications of the deception story first. If the intended victim has already formed an estimate of the most likely course of action, the practitioner need only take those actions necessary to provide substantiating evidence. Once the victim has focused on a single most likely course of action, receipt of later information will be evaluated in terms of whether or not it matches the current hypothesis. The victim may then ignore contradictory evidence, fit ambiguous evidence to match the hypothesis as if no ambiguity existed, and accept deceptive activity with little scrutiny.

Cover and Deception and the Intelligence Process

The use of C & D extends across the conflict spectrum, and applies to other dimensions of military/political analysis. Cover and deception has been successfully applied

The principle objective of the intelligence process is knowledge the opposing decision-maker's plan. Usually, this cannot be gained directly. Further, the raw data which the intelligence system receives is limited both by the attributes of military activity which are observable and by the time slices when collectors are actually tasked to collect. Finally, the interpretation of activity depends on the accuracy of our concept of the opponent's force procedures. Even without any intent to deceive by an opponent, the limitations of the intelligence process would leave us with an incomplete and sometimes misleading picture of his activities and objectives.

The threat assessment process is vulnerable to C & D at each step. Figure 2 decomposes the process, and shows the opportunities for a skillful opponent to employ C & D. To begin with, an adversary can control the timing and type of activity by his forces to manage our perception of the observable features of his course of action. Assuming that an analyst began with an accurate baseline of enemy locations and activity, this type of deception would cause errors in threat situation monitoring-monitoring the current state of enemy forces. At this point, the opponent loses direct control over his ability to manage perception, but must rely on the weaknesses of our intelligence analysis system. The opponent's C & D plan attempts to orchestrate observable activity so that collection distortions and interpretation errors are propagated through the higher levels of intelligence analysis.

Indications analysis, relying on an incorrect statement of the current situation, will misdirect requests for additional collection. Key activities will be missed, and others will be assessed as having occurred when they have only been simulated. The final step in the process, threat synthesis, matches key activities with the hypothesized set of courses of actions. If those key activities are not correctly identified, or if the set of courses of action is incomplete, an incorrect assessment of the opposing decision-maker's intentions will be given to our own commanders.

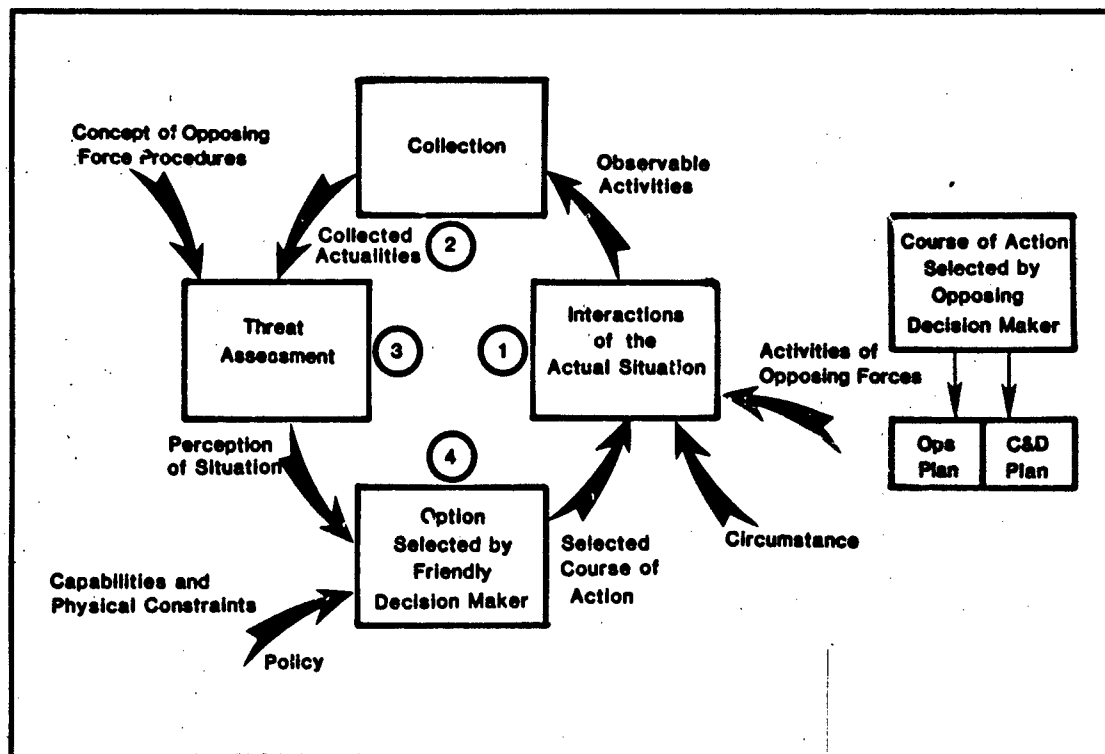


FIGURE 1 OPERATIONAL CONTEXT FOR C&D

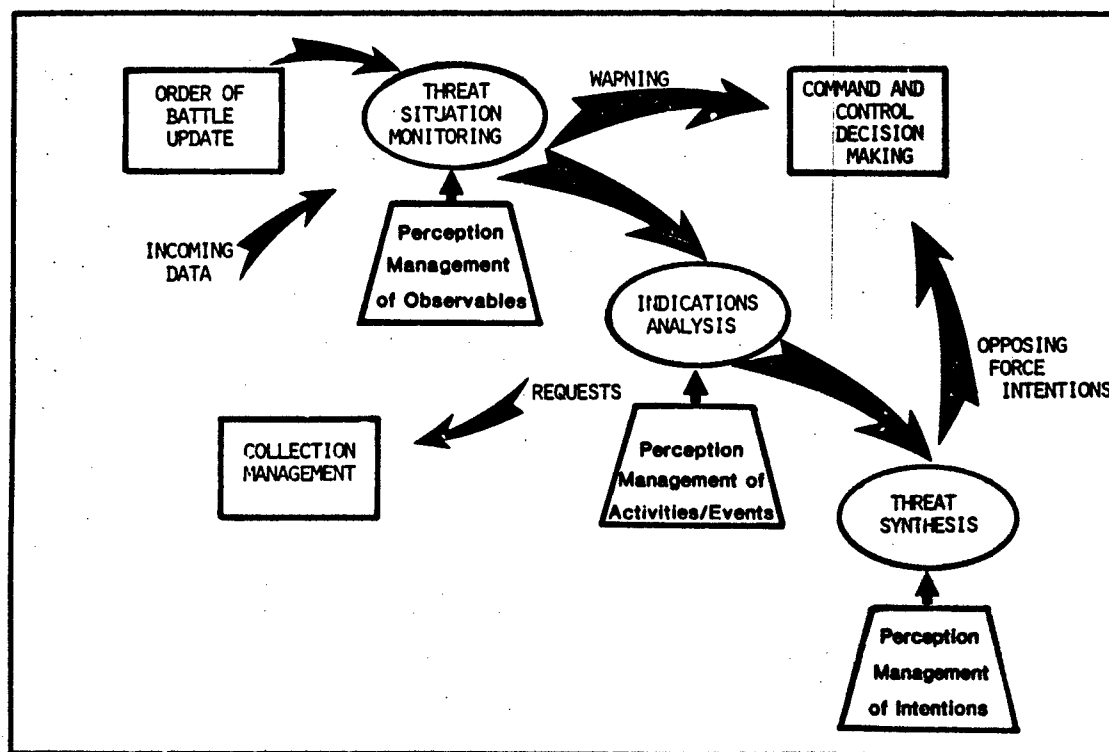


FIGURE 2 THREAT ASSESSMENT VULNERABILITIES TO C&D

Unfortunately, many of the developments in intelligence systems in recent years have increased our vulnerability to C & D at the same time that they have increased our ability to collect data and to analyze activity. Together with a greatly increased ability to collect, we have developed systems to help the human analyst exploit that capability by focusing attention on items which analysts have identified as keys.[2] Thus, each enemy course of action can be broken down into indicators-steps which must be taken to realize that action, indicators into key activities, activities into observables. The result is a system of great power for focusing attention on significant pieces of information and for leading to conclusions of intent based upon a clear path of reasoning. The weakness of this system is that the discriminators at each step become high value targets for an opponent's C & D activities, and as we discuss in the next section, it is highly likely that the Soviets will attempt to exploit this weakness.

Soviet Cover and Deception Doctrine and Applications

Extensive use of cover and deception techniques in the tactical environment are basic tenets of Soviet military doctrine. Natural Soviet proclivity to secretiveness coupled with Soviet experience and lessons learned over the last thirty years have convinced Soviet leaders that cover and deception are invaluable tools in tactical warfare. These attitudes have undoubtedly been strengthened by the successful application of C&D by both Arabs and Israelis in the Middle East wars, and by the British in the Falklands.

The commonly held Western appreciations of Soviet C&D capabilities may already be out of date. The Soviets have repeatedly demonstrated their ability to make the improvements necessary to bring capability in other fields up to the demands of doctrine. A clear example of this is the comparatively recent development and mass deployment of equipment (such as the KIROV-class VSTOL carrier) which make the extension of Soviet offensive doctrine into the naval domain a credible threat to NATO. A corresponding effort to increase the level and sophistication of their C&D capabilities can therefore be hypothesized as a most likely Soviet course of action.

Soviet Cover and Deception Doctrine

Soviet doctrine for Cover and Deception derives from the doctrinal requirement for surprise, one of the basic principles of what the Soviets call "operational art". These basic principles, in the order assigned by the Soviet author Savkin [3], include:

- Mobility/Tempo
- Concentration of Efforts
- Surprise
- Activeness of Combat
- Preservation of Effectiveness
- Conformity of Goals to Conditions
- Interworking

Mobility/Tempo includes not only speed of movement but also flexibility, such as in changing the axis of an attack. Concentration of efforts is more familiar to us as the principles of mass and economy of force. By "Surprise", the Soviets mean the ability to force an enemy to fight in a situation unfavorable to him- either in a place or time which does not allow him to make full use of his own forces. "Activeness of combat" states the Soviet desire to hold the initiative; this is also the principle of the offensive. "Preservation of Effectiveness" is the Soviet reaction to the advent of weapons of mass destruction, which require that their forces avoid premature concentration, and that they be equipped to survive in a CBR environment. "Conformity of Goals to Conditions" demands that the commander assign reasonable goals to his forces. Because commanders adhere to this principle, the subordinates can therefore be held

accountable for any failure. "Interworking" is the basis for the Soviet combined arms approach to military operations. Interworking refers not only to the coordination leading to joint efforts by combat forces, but also to the coordination of front and reserve units, combat and support units, center and flanks.

One of the significant features of these principles is the degree to which each is often implemented in terms of the others. A surprised enemy, for example, is given no time to recover if the attacking forces maintain the tempo of their attack, and keep the initiative. Also, the dispersion prior to attack necessary to preserve combat effectiveness makes it more difficult for the enemy to determine the time and place of an attack. Finally, the Soviets achieve consistency in their deception plans by a combined arms approach to C&D operations.

In the discussion of methods to achieve surprise, Savkin [4], mentions six general types:

- Lead the enemy astray
- Secrecy or Preparation
- Unexpected Use of Nuclear Weapons
- Deliver attacks at Unexpected place/time
- New means/methods of warfare
- Avoid repetition of methods

The first of these methods, leading the enemy astray with regard to one's intended course of action, is the doctrinal basis for Soviet deception operations, just as secrecy of preparation is the basis for the widespread Soviet use of cover and camouflage for offensive purposes. The last two methods are useful in understanding how the Soviets have been able to continue to surprise their opponents in intervention actions over the past thirty years. In his discussion of "new means and methods of warfare", Savkin explains that this is usually achieved by using existing means in ways unknown to an enemy, rather than by the introduction of a totally new capability. This, together with the avoidance of repetition in the methods of operations, including deception operations, put our intelligence system on notice that hypotheses limited to past patterns of Soviet actions not only fail as aids to detection of new patterns, but also increase the probability that the Soviets will exploit our tendency to correlate the elements of a new course of action with an old one.

Soviet C&D Experience

A brief summarization of four Soviet C&D operations, beginning with the successful preparations for Operation Bagration in the summer of 1944, [5] illustrates how far the Soviets have progressed in the ability to lead their opponents astray. Knowing the German preoccupation with defending their economic base, in this case the Ukraine, the Soviets covered their preparations for an attack in Belorussia through denying the Germans any information that would contradict that hypothesis. For this reason, the Soviets moved their forces to their jump-off positions under the cover of darkness, and spread the observable indicators of impending attack over the entire eastern front- aerial reconnaissance, bomber sorties and air defense were not concentrated in the central front. Communications activity by units dedicated to the attack was kept to a minimum. Although the Germans expected a summer offensive, Soviet security precautions together with German assessment that the most likely location for an attack was the area which they most feared to lose, combined to leave the Germans unprepared for the Soviet assault. A significant difference between the Soviet C&D operations and the coincident US/British effort to deceive the Germans as to the location of the cross-channel landing, however, was that the deception ended when the attack began. This difference has continued in Soviet operations to this day.

In their preparations for intervention in Czechoslovakia in 1968, the Soviets exploited both Western and Czech preconceptions of the sequence of events which would precede such an action. In 1956, the Soviets had called the Hungarian leaders to Moscow and then invaded. In 1968, the Soviets moved their forces to the military districts bordering Czechoslovakia for a long series of exercises, summoned the Czech leaders to Moscow, and did not invade. Instead, the Czechs were allowed to return home believing the Soviets would withhold action as long as the Czechs committed no excesses of liberalization. Some Soviet units were recalled from the border areas. Although the Soviets still had sufficient force in place to intervene, both Czechs and the West changed their assessment of Soviet intentions. Both were therefore unprepared when the Soviets, with token elements of other Warsaw Pact nations, invaded in August. It is likely that at least part of the reason for the timing of the invasion was to coincide with the summer vacation season in Europe, when many European decision makers would be on vacation. In addition, the Soviets used a ruse to gain control of the main airfield outside of Prague, sending the landing control party on an Aeroflot flight dressed as tourists. The "tourists" easily overpowered the Czech control tower personnel and then proceeded to handle the landing of the aircraft carrying the leading elements of the invasion force.

The Soviet invasion of Afghanistan [6] is a good illustration of how the Soviets were able to achieve the same end-control of the capital city airfield - while varying the means. This time the Soviets flew in the airfield control party weeks before the invasion as reinforcements for Soviet units already deployed there since September, thus arousing no curiosity. The Soviets then disarmed the Afghani armored forces by recalling the Afghan ammunition and anti-tank guns for inventory, some of their tanks for winterization and others for the repair of defects. Although western intelligence was not surprised - the US had warned the Soviets twice against intervening in Afghanistan - the Soviet choice of Christmas as the invasion date meant that any Western reaction would be delayed as the leaders hurried back from their vacations. The surprise achieved against the Amin government is reflected by the ease with which Soviet divisions occupied the country.

In its initial stages, the Soviet reaction to the internal political events in Poland were similar to those of 1968. [7] A long series of large-scale military exercises were held in the western military districts, the Polish premier visited Moscow, and no immediate invasion occurred. US and NATO intelligence recognized the clear threat of intervention behind the exercises, and exerted diplomatic pressure upon the Soviets to allow the Poles to settle the party-union conflict. The nature and timing of the coup in December therefore caught both Solidarity and the US and its Allies by surprise. Again the Soviets chose a time (Sunday morning during the Christmas season) when opposing decision makers would be at home, and an unexpected means - the Polish internal security forces. The intended victims of the deception were successfully deceived even though quite sensitive to the possibility of deception. The exercises were widely perceived as a cover for the Soviets' true course of action, but this knowledge did not result in correct assessment of Soviet intentions.

In each of the cases sketched above, a great deal of information indicating the true Soviet courses of action was available. The tactical observables of the Soviet preparations were collected. The capabilities of the forces involved were known as was the general intent. In each case, however, the Soviets succeeded in deceiving their victims as to timing and method. This series of successes makes the development of C&D countering techniques, discussed in the next section, a necessary part of any effort to improve the performance of our own intelligence system.

Developing Techniques to Counter Cover and Deception

Countering deception is a two-step process. The first is to identify the targets for deception, the second is to identify how C&D directed against those targets can be recognized and then exploited. These techniques must then be integrated into both sensor related improvements and into the development of expert systems and other ADP tools for intelligence analysis.

Identifying C&D Vulnerabilities

Identifying C&D vulnerabilities can benefit from the great effort already expended in the development of structured indications and warning systems. These systems break down a range of courses of action into the steps (indicators) required to achieve them, decompose these steps into their key activities, and then identify the observables associated with each key activity. These observables are the high value targets for C&D operations, since by managing an opponent's collection of these observables, the deceiver exerts control on the basis of the victim's perceptions. In order to understand how perceptions can be managed, it is therefore necessary to begin by identifying the sources and methods used to gather intelligence data (see Figure 3). Identifying the targets of

COLLECTION DISCIPLINES SENSORS											
	ACINT	COMINT	ELINT	HUMINT	IRINT	MAGINT	OPTINT	PHOTINT	RADINT	TELINT	VISINT
AIRBORNE		●	●		●	●	●	●		●	●
AIRBORNE DETECTION									●		
AIRBORNE IMAGING									●		
ATTACHE				●							
CASUAL				●							
DEDICATED				●							
FIXED SITE ACTIVE	●										
FIXED SITE PASSIVE	●										
MOBILE ACTIVE	●										
MOBILE PASSIVE	●										
OVERHEAD		●	●		●	●	●	●		●	●
OVERHEAD DETECTION									●		
OVERHEAD IMAGING									●		
SUBSURFACE		●									●
SURFACE			●		●		●	●		●	●
SURFACE FIXED SITE	●								●		
SURFACE MOBILE SITE	●								●		

FIGURE 3 COLLECTION DISCIPLINES AND SENSORS

C&D involves a "reverse engineering" process (see Figure 4). We know how individual collection sources and disciplines can be exploited by an opponent in a general way, and we know which disciplines are employed to collect given observable. By matching collection means and C&D methods pairs to the association of collection means and observables, we can construct a C&D matrix for each key activity. Figure 5 presents a sample C&D means matrix for one possible key activity - the deployment forward of a technical unit (such as a bridging unit). Large scale deployment of such units would be required prior to the initiation of hostilities. Each row in the matrix summarizes how the observable within a particular collection discipline




FIGURE 4 DEVELOPING TECHNIQUES TO COUNTER C&D: IDENTIFYING VULNERABILITY

ACTIVITY TO BE CONCEALED OR SIMULATED: DEPLOY TECHNICAL UNIT FORWARD	
COLLECTION DISCIPLINE	COVER & DECEPTION MEANS
ACINT	•SIMULATE SOUND OF VEHICLES
COMINT	•MAINTAIN NORMAL LEVELS OF COMMUNICATIONS AT GARRISON LOCATIONS TO MASK MOVEMENT OF UNITS •MAINTAIN STRICT COMMUNICATIONS SILENCE BY THE MOVING UNITS
ELINT	•STAGGER RADAR CHECK-OUT BEFORE DEPLOYMENT TO SIMULATE NORMAL ACTIVITY •MAINTAIN NORMAL LEVELS AND TYPES OF RADAR ACTIVITY IN GARRISON AREAS
HUMINT	•RELEASE COVERING EXPLANATION FOR ACTIVITY (E.G., EXERCISE ANNOUNCEMENT, TROOP ROTATION)
IRINT	•DEPLOY THROUGH AREAS WITH HIGH LEVELS OF BACKGROUND HEAT (URBAN AREAS, MAJOR ROADS) •DEPLOY THROUGH AREAS WHICH ABSORB IR EMISSIONS •SIMULATE NORMAL GARRISON ACTIVITY WITH NONESSENTIAL VEHICLES
CPINT	•CAMOUFLAGE DEPLOYING VEHICLES AS NON-MILITARY •SIMULATE ESSENTIAL VEHICLES IN GARRISON WITH NONESSENTIAL ONES •MOVE AT NIGHT
PHOTINT	•SIMULATE ESSENTIAL VEHICLES WITH DUMMIES IN GARRISONS
VISINT	•DIVERT FOREIGN OBSERVERS FROM DEPLOYMENT ROUTES •ALLOW OBSERVERS TO SEE STAGED ACTIVITIES ELSEWHERE

FIGURE 5 EXPANSION OF SAMPLE MEANS MATRIX

could be simulated. The more elaborate the deception, the greater number of these methods would be employed, and the larger the number of units simulated.

Uncovering C&D depends upon discovery of inconsistency. The Soviets are likely to apply their military doctrine to C&D operations, in that they will strive to achieve consistency with the least effort necessary (economy of force), and with integration of C&D operations in all domains (interworking). To counter C&D therefore means to progress from line items in the means matrix to correlation of line items in the matrix, and then to the higher levels of the indications structure, as well as to other intelligence disciplines (for example, the order of battle). This search for inconsistency takes place on three levels, each demanding a higher level of man (or machine) intelligence:

- Single collection discipline
- Multiple collection disciplines
- Analytical procedures involving one or more intelligence disciplines

Each technique, properly employed, should stretch the deceiver's web of consistency in the observables harder, until finally it gives way.

Single Discipline Techniques

Single discipline techniques address the weaknesses in the collection and interpretation chain, and can be divided into two types;

- bringing the target into the field of view and
- increasing target discrimination.

Success in either of these can be achieved by improving either the sensor capabilities, the exploitation processing, or even by alerting interpreters to the likelihood of a particular C&D method.

As an example of the application of these techniques, consider an enemy attempting to deploy SAM units forward, and attempting to cover the radar emissions. He may try to keep them out of our ELINT field of view by restricting the time and power of his emissions. They would be brought into the field of view by expanding the duration of coverage, or by deploying more sensitive sensors. This affects both the enemy's ability to avoid coverage, and his ability to remain undetected in the presence of a collector. The enemy might also seek to cover the mass deployment of radars forward by testing them individually, so that the overall level of SAM radar activity in a given area does not change. This can be uncovered by increasing the ability to discriminate among the radar signals of different units with the same types of radars.

Multiple Discipline Techniques

Multiple discipline techniques seek to break down inconsistencies between two or more observables associated with the same key activity. The first step in applying these techniques is to take advantage of the means matrix to identify the opportunities for multiple discipline correlation. There follows a determination of whether the current collection schedules for the sensors involved allow simultaneous coverage. Planning for such coverage increases the burden of activity required to maintain a deception. Finally, the analyst's ability to make effective use of multi-source coverage requires that the interpretation of the collection be organized by activity. For the radar in the above example, therefore, the analyst would be given the PHOTINT, ELINT and other coverage for a particular area over a specified time range. Multiple discipline techniques make cover difficult at the tactical level, and make simulation extremely difficult, since the effort required to simulate an activity in many different

disciplines may be greater than the effort required by the activity itself.

Analytical Techniques

Beyond current collection and exploitation the intelligence analyst can uncover a C&D operation by comparing current activity with the knowledge base of an opponents capabilities and options. These comparisons are intelligence cross-discipline consistency checks, in which current intelligence is matched with basic intelligence on the one hand and threat assessment on the other.

Basic intelligence provides the analyst with a reference of what an opponent can do. This includes the physical capabilities of equipment—can a mobile radar deploy from A to B in a given time? In addition, it provides an organizational and doctrinal reference for current activity. These are particularly useful in evaluating the activity of the Soviet military, which has minimized organizational variations and which does not encourage deviations from standard operating procedures. A simulated SAM battalion, for example, must include the correct number and relative location of radars, launchers and communications equipment. From the organizational and doctrinal standpoint, it must be co-located with one of a limited number of other types of units. Discrepancies in any of these factors becomes the basis for requests for additional collection, and for expanding the scope of the analytical evaluation.

Once the time and space relationships between indicators and an opponent's likely courses of action have been established, these can also help uncover a C&D operation, and can also help the analyst recognize when the actual course of action does not match any of the hypothesis. One of the major benefits of the structured warning systems is that the analyst can be alerted to the inconsistent absence of activity. This absence could occur under any of the following conditions:

- the activity is present, but is being covered
- the activity not present, and the other key activities are being staged
- the activity is not present, and the other activities are part of a course of action outside the current range of hypotheses

The analyst can identify which explanation applies by increasing collection and exploitation effort to uncover activities if it exists. If it is not found, solutions must be sought along both collection and analysis paths.

In the collection/interpretation domain, increased effort would be applied to determine if some of the observed activities are actually simulations. At the same time, the threat assessment process needs to reevaluate whether the absent activity is a necessary part of a course of action, and whether a new hypothesis would be consistent with the current combination of active and inactive indicators. The discovery of C&D operations during this process has an additional value in that their existence is itself an indication of an opponent's course of action.

Conclusions

Current collection, exploitation and intelligence analysis systems are vulnerable to cover and deception. This vulnerability has increased as analytical aids have tended to focus on the observables associated with the key steps for a limited range of courses of action. Soviet doctrine, with its emphasis on surprise and continuing variations in the means of realizing surprise, is ideally suited to exploit those limitations, and they have been uniformly successful in applying this doctrine to the present day. Both current analysis aids, which are essentially

production systems, and the expert systems now under development [8] are basically similar to commercial systems developed for medical, geological or engineering applications. [9] For successful application to military intelligence, the technology must be "hardened" to withstand the skillful use of Cover and Deception.

References

- [1] Richards J. Heuer, Jr., "Strategic Deception: A Psychological Perspective," presented at the 21st Annual Convention of the International Studies Association, Los Angeles, California, March 19-22, 1986.
- [2] Ralph F. Gerenz, "A Methodology for Improving the Strategic Warning Process," Journal for Defense Research, Crisis Management Edition, April 1977.
- [3] V. Ye. Saukin, The Basic Principles of Operational Art and Tactics (A Soviet View), Moscow, 1972. USAF Translation, Chapter 3, pp. 167-277.
- [4] *ibid.*, pp. 230-239.
- [5] Barton Whaley, Strategem: Deception and Surprise in War, published in manuscript by M.I.T., Cambridge, Massachusetts, 1969, pp. 144-151.
- [6] Jiri Valenti, "Perspectives on Soviet Intervention-Soviet Use of Surprise and Deception," Survival, March/April, 1982, pp. 55-56.
- [7] *ibid.* pp. 57-60.
- [8] David Brown and Harvey S. Goodman, "Artificial Intelligence Applied to C3I," Signal, March, 1983, pp. 27-33.
- [9] Avron Barr and Edward Feigenbaum, The Handbook of Artificial Intelligence, Volume 1, pp. 190-199, Los Altos, California: William Kaufman, 1981.

DISTRIBUTED ESTIMATION IN THE MIT/LL DSN TEST-BED*

J. R. Delaney, R. T. Lacoss and P. E. Green

Massachusetts Institute of Technology
 Lincoln Laboratory
 Lexington, MA 02173

I. Introduction and Summary

A DSN is a surveillance and tracking system employing many geographically dispersed sensor/processor nodes connected by a computer communications network and implemented as a confederacy of identical autonomous cooperating processes. The general properties of a DSN are being investigated through the development and exercise of a test-bed system. The detection and tracking of low-flying aircraft using simple acoustic sensors has been selected as a specific problem to be addressed. Each node in the test-bed consists of an array of microphones as the sensor, several small computers for data processing, and a digital radio for communicating information between the nodes. Use of microphone arrays as sensors limits each node to measuring target azimuths; nodes must exchange target azimuth information in order to track target positions.

A DSN need not be as simple or as homogeneous as the test-bed. It can mix sensor types such as radar and passive IR, putting one type or both at any node. Data processors need not be colocated with all sensors and vice versa. Data processing capacity can vary from node to node and communications capacity can vary from link to link. The communications network can mix broadcast radio, point-to-point radio, and wire. But it is our opinion that all of the significant DSN-related technical issues which must be addressed in developing such complex systems must also be addressed in developing the test-bed system and that the relative simplicity of the test-bed helps focus attention on those DSN-related issues.

This paper describes the acoustic tracking algorithms currently used in the MIT Lincoln Laboratory (MIT/LL) Distributed Sensor Networks (DSN) test-bed. It discusses the original motivation for inclusion of various features in those algorithms and the lessons learned about those features through experimentation with real and simulated data. Plans for modifications to the detection and tracking algorithms are briefly sketched.

Acoustic propagation introduces delays in azimuth measurements which complicate the merger of target azimuth information into target position estimates. At any one measurement time, each sensor measures target azimuths that correspond to different times in the past, depending on target-to-sensor range. A node cannot estimate target position from just the most recent measurements because they correspond to different target positions. Each node must maintain a history of past sensor measurements in order to track targets. Two algorithms have been developed for estimating target positions from histories of sensor measurements produced by two or more nodes.

One algorithm has the advantage that it introduces no more delay into the target detection and tracking process than is forced by acoustic propagation. But at one measurement time, it usually estimates the positions for different targets at different times. Over a regular sequence of measurement times, it usually estimates the position of the same target at irregularly spaced times.

The other algorithm forms estimates at each measurement time with a specified delay which is uniform for all targets. The product of the delay and the speed of sound is the maximum range at which targets can be tracked, so that delay is usually set equal to the sensor detection range divided by the speed of sound. On the surface, the generally larger delay of the latter algorithm makes it appear undesirable. However, the irregularity of estimation times produced by the first algorithm so complicates further data processing as to make it even less attractive.

The microphone arrays used in the test-bed have a detection range on the order of 5 to 10 km and the radios have a similar range. Thus, no one node can measure a target's azimuth for very long. This difficulty can be overcome if each node's azimuth data is distributed to every other node. But to do so would be difficult in a large DSN, producing heavy communications traffic and requiring much redundant data processing in each node. The least distribution practical, broadcast of azimuth data only to nodes in direct radio contact, is done in the test-bed so as to hold down communications traffic.

Only azimuth data is now exchanged by the test-bed nodes. Because this data is exchanged only between nodes in direct radio contact and because the radio range is comparable to the sensor range, each node has sufficient information to track the positions of targets within its sensor coverage. This restriction impacts target position tracking by individual nodes only in the area of track initiation. As a target moves through the DSN, each node must acquire the target as it enters the node's sensor coverage and initiate a new position track for the target despite the likely existence of position tracks for that target in other nodes. This restriction also has a major impact on the surveillance function of the DSN as a whole because individual nodes have a myopic view of the targets within the coverage of all the DSN's sensors.

Intuitively, exchanging position tracks as well as azimuth data between nodes in direct radio contact should allow target position tracks to be handed over instead of reinitiated. And more extensive communication of position tracks should allow formation of a complete surveillance picture. But it is not yet clear how to combine, in a statistically

valid manner, position tracks formed in different nodes using sensor data which is not all exchanged. Recent research has examined related, simpler problems, and suggests that this difficulty is surmountable. However, the applicability of the research results needs further study.

II. MIT/LL DSN Test-Bed

Figure 1 shows a simple block diagram of one DSN test-bed node as well as photographs of a microphone array and a mobile node vehicle carrying the associated computers, radio, and their power supply. Atmospheric conditions limit the maximum range at which typical targets can be detected to between 5 and 10 km. The test-bed radios, currently under construction, will typically be limited to a similar range by line-of-sight propagation. Thus, each node communicates directly with those nodes having overlapping sensor coverage and only those nodes. The radios are also designed to measure the range between nodes, allowing the test-bed to estimate the relative locations of its nodes as indicated in the figure. Until the packet radios are available, broadcasts are simulated using wire communications.

Figure 2 expands the tracking block in Figure 1 and includes, for reference, the microphone array and the signal processor. It shows all of the data flows in the detection and tracking system, including those between nodes. The data flows are complex but, as will be seen, the complex flows are necessary. The data from each microphone array is processed every two seconds using adaptive, nonlinear filtering¹⁻³ to detect local maxima of the incident sound power (averaged over the preceding two seconds) in frequency and azimuth angle. A measurement of each detection's frequency, azimuth angle, and average sound pressure level is produced in the process.

The detections and measurements are made over a four to five octave frequency band. Such broad bandwidth allows detection of multiple harmonic emissions by a single target. Sensor data conditioning is done to reduce tracking computational load by clustering together detections which could plausibly have been caused by a single target and by discarding detections which are relatively weak in sound pressure level. Each cluster is characterized by the average azimuth angle of the component detections (weighted by sound pressure level) and by the total sound pressure level. The average azimuth angles are treated as target measurements thereafter.

Acoustic propagation variability makes incident sound pressure level a second-order measure of target range; the sensors are primarily azimuth-only measurement devices. Thus, the full target state is not observable using a single sensor and no node can accurately estimate it without information from other nodes; the nodes must share measurement information to form an estimate of the full target state.

Low-flying aircraft can travel at an appreciable fraction of the speed of sound and even exceed it in some cases. As a result, the sensors measure target azimuth angle not at the measurement times but at the times when the measured sounds were emitted. Target azimuth angles corresponding to present target positions are not observable. The lack of observability prevents the direct application of familiar tracking techniques, including the Kalman filter, which update target state estimates using current measurement data only. Further complications are the nonlinearity of the measurement process and

its effective noncausality; a measurement by one sensor may be made later than a measurement by another sensor but may contain information about the target state at an earlier time in the target's frame of reference.

Target azimuth data must be accumulated over time in order to form target position tracks. The first step in this accumulation is target azimuth tracking. The azimuth tracker in each node is quite conventional, with the exception that it does not estimate the unobservable present target azimuth angle but rather the observable "acoustic" azimuth angle. The tracker uses a two state (azimuth angle and azimuth angle rate) α - β tracking algorithm for estimation and prediction. Data association is done quite simply. A measurement is associated with an azimuth track and used to update the state estimate if the azimuth angle lies within a window about the filter's azimuth angle prediction for the measurement time. Only one association is allowed per track or per detection. Should a measurement fail to associate with any existing track, it is used to initiate a new track. A newly initiated track is terminated if no measurement is associated with it at the next measurement time; any other track is allowed only one missed data association in a row.

Full azimuth track state estimates are maintained within each node for the targets sensed locally but only the azimuth component is broadcast to other nodes. The azimuth components computed locally are also added to the node's azimuth history data base along with all azimuth components received from other nodes. Each azimuth track created by a node is given a unique (within the node) identifier which is broadcast with each azimuth component. This tagging allows azimuth components broadcast at different times but based on the same azimuth track to be associated. Figure 3 sketches the organization of the azimuth history data base. It is tree-structured, with the data sorted by originating node and azimuth track.

Only the data in the azimuth history data base issued in further processing. The effect is to treat the each node's sensor, signal processor, sensor conditioning, and azimuth tracker as a virtual sensor with the measurement properties of the whole chain. Each node can be thought of as connected to a number of such virtual sensors, one at its location and the others at the locations of those nodes with which it has direct radio contact. This artifice helps modularize the tracking system, decoupling details of the sensor, signal processor, etc., from the remainder of the tracking system.

The remainder of the tracking system must form estimates of target dynamic state (position and velocity) from these virtual measurements of acoustic azimuth. The process is done in two steps at each node:

- 1) Estimate target positions from acoustic azimuth data in the azimuth history data base, and
- 2) Estimate target dynamic states from the estimated target positions.

Azimuth histories can be combined in two ways to produce position estimates. The current tracking system uses the reflection algorithm^{4,5}. The algorithm takes as inputs two azimuth histories originating from different nodes and the locations of those nodes. If the two azimuth histories could plausibly be associated with a single target, the algorithm estimates the position of that target for the

emission time of the sound just detected at the node closer to the target. A position estimate produced by the algorithm can be regarded as a "measurement" of the target's position, delayed in availability depending on the target's range from the sensor. The delay is the least achievable with any position estimation algorithm; it equals the shortest propagation delay in the current azimuth measurement data.

In each node, the reflection algorithm is applied to all legitimate combinations of azimuth histories. The resulting position estimates for different targets are for different times. The times can range from the most recent measurement time (for targets overflying sensors) to 30 seconds earlier (for targets 10 km. from the nearest sensor). Because target-to-node range varies with time, so does the delay for a single target's position estimates at sequential measurement times, even if the estimates are based on the same pair of azimuth histories (same nodes and track identifiers) at sequential measurement times.

The variable times associated with target position estimates produced by the reflection algorithm complicate the process of building position tracks from position estimates. Consider a target which is tracked simultaneously by three test-bed nodes and assume that each node is within radio range of the other two. Then each node contains three azimuth histories corresponding to the target. Applying the reflection algorithm to each pair of histories would typically yield three different position estimates, each for a different time. It is very difficult to recognize that all of these position estimates correspond to the same target, even if the azimuth histories are completely accurate and if the reflection algorithm introduces no inaccuracies.

For this reason, position tracks are currently formed only from position estimates based on particular pairs of azimuth histories, i.e., for the same pair of nodes and azimuth tracks within those nodes. To facilitate this isolation of position tracks, they are kept in a tree-structured data base organized like the azimuth history data base. In this case, entries are sorted first by the originating pair of nodes and then by the pair of identifiers of the originating azimuth tracks (see Figure 4). Position track data base entries consist of a time and estimates of target easting, northing, east velocity, and north velocity at that time, plus some auxiliary information.

A position track is updated when a new position estimate is produced for the pair of azimuth tracks histories upon which that track is based. A new position track is created if no entry exists in the data base corresponding to an azimuth track pair which passed the reflection algorithm's test. A α - β tracking algorithm is again used for prediction and estimation. Azimuth histories are not completely accurate and the reflection algorithm can amplify those inaccuracies. The resulting position estimates can be inaccurate not only in position but also in time. For this reason, position estimate times are smoothed before the α - β algorithm is applied. The auxiliary information in each position track data base entry is the state of the smoother for that position track.

Neither target position estimates nor position tracks are exchanged between nodes in the current test-bed. If the nodes are all directly connected by radio to each other such an exchange would be

unnecessary. Each node would (in the absence of communications failures) have identical azimuth history data bases and produce identical target position estimates and position tracks. But complete connectivity is not practical in large DSNs; thus, azimuth history data bases typically differ from node to node and so must the position estimates and position tracks based upon the azimuth history data. Each node then has an incomplete picture of the targets in the DSN's coverage. This point is illustrated in the next section.

Exchanging target position estimates or position tracks could provide nodes with additional information, but consider an extreme example of what can go wrong if target state estimates are not properly combined:

A local state estimate α is created in data processor A from a single measurement by one sensor. That estimate is shared with data processor B, which uses it to create an identical local state estimate, β . Data processor B shares its local state estimates with data processor A, including β . Data processor A associates local state estimates α and β with the same target and combines them to produce local state estimate λ with half the variance of α and β . This process could ultimately lead to the existence in data processors A and B of a local state estimate ω of infinitesimal variance based on a single sensor measurement.

Given the complexity of this issue, the decision was made to limit the initial version of the test-bed to exchanging only the azimuth components of azimuth tracks and only between nodes in direct radio contact. Later, versions of the test-bed will experiment with more extensive data exchanges.

III. Tracking Performance

All of the results shown in this section are for one field experiment^{2,3,6}. Four test-bed nodes were used to record the passage of a UH-1 helicopter west-to-east along the flight path illustrated in Figure 5 at a ground speed of roughly 65 knots and at roughly 1000 feet above ground level. The letters F, H, J, and L indicate the locations of the four nodes. The circles indicate checkpoints used by the helicopter pilot to maintain his flight path and by observers to time the helicopter's passage.

The runways overflown near checkpoints 3 and 4 are Hanscom Field, a busy military and civilian airfield. Under checkpoint 6 is Route 128, a heavily used eight-lane superhighway. Another major highway, Route 2, lies just south of the area shown on the map. Normal activity at Hanscom field provided additional aircraft in sensor coverage during the experiment and normal activity on Route 128 provided an extended and irregular interfering sound source. Construction vehicles and stationary mechanical equipment operating near the microphone arrays also provided acoustic interference.

Figure 6 is an azimuth-time intensity plot for the output of the signal processor at node F. The plot only includes those measurements at each time having a sound pressure level within 10 dB. of the maximum value at that time. The curves overlying the measured values are the sequences of azimuth components of the azimuth tracks produced when processing the data. Each continuous curve corresponds to one track. Because of breaks in the measurement data, several distinct tracks (having distinct track identifiers) are caused by each sound source.

The curve marked "TRACK" is the track of the UH-1 helicopter, which reached its point of closest approach at roughly 240 seconds and was roughly 600 meters from the microphone array at that time. Two other sound sources were tracked: a fixed-wing aircraft with a varying azimuth angle and a bulldozer with an essentially constant azimuth angle (due south). The "speckles" on the plot are signal processing artifacts or very intermittent sound sources. Such artifacts or sources are detected uniformly over time, but are suppressed by the sensor data conditioning process early in the plotted data because of a relatively loud target near the sensor.

Figure 7 shows the position tracks formed using updated azimuth estimates calculated in node F (shown in the previous figure) and in node H. Each cross is the position component of a newly updated position track. Thin lines connect sequential position components of individual position tracks. The line marked "TRACK" is again the track of the UH-1 helicopter. That track is crossed by another, that of the fixed-wing aircraft. The other tracks are of intermittent sound sources or are processing artifacts, e.g., "ghost" tracks based on erroneous pairings of azimuth track histories which manage to pass the reflection algorithm test. The short, dense track radiating from node H is a "ghost". Ghost tracks can often be recognized because they trace out physically unreasonable trajectories, beginning or ending at a node or involving unrealistic accelerations.

Recording of the measurement data allowed experimentation with different connectivities between the nodes. Figure 8 shows the position tracks formed in node J when it received azimuth data broadcast only by nodes H and L. The figure includes distinct but overlapping tracks of the UH-1 helicopter, each derived from a different pair of azimuth histories. Note that node J has an incomplete surveillance picture. It does not include the track of the fixed-wing aircraft which appears in Figure 7. The latter track was based in part on azimuth information sensed at node F; information unavailable to node J in this case.

IV. Lessons and Plans

Use of the reflection algorithm for target position estimation has the advantage of producing the most up-to-date position estimates possible. But it has the disadvantage of producing position estimates with time-varying delays, making it difficult to recognize position estimates associated with the same target. A consequence is the overlapping UH-1 position tracks in Figure 8. Location estimate time smoothing was required in the position tracking algorithm as a consequence of inaccuracies in the time-varying delays. Experiments using simulated measurement data with no azimuth measurement inaccuracies has revealed that this smoothing process by itself produces some inaccuracies in the position tracks. These observations have led us to question whether the advantage of the reflection algorithm is worth the disadvantage.

Another position estimation algorithm, the possible position algorithm^{5,7,8} takes the same inputs as the reflection algorithm and produces position estimates with a fixed delay. Targets at ranges great enough for acoustic propagation time to exceed that delay are ignored by this algorithm even if they are detected. So the delay is usually chosen to equal the propagation time for sound from a target at the maximum detection range. This delay would be

30 seconds for the test-bed. For a target very near a node, the possible position algorithm cause a delay nearly as large between detection of the target by that node's sensor and the creation of a corresponding position estimate. Lesser delays would occur for targets further from nodes. Since targets are typically first detected away from all nodes and are well in track by the time they are close to any node, this disadvantage is probably less significant than it appears on the surface.

The availability of position estimates are regular and common times should allow significant simplification of the position tracker. Location estimates plausibly caused by the same target could be clustered together in the same manner as are azimuth measurements produced by the signal processor. Entries in the position track data base would not need to be sorted by originating azimuth history pairs and more usual data association could be done. Smoothing of the position estimate times would be unnecessary, eliminating this potential source of error. Experiments using the possible position algorithm for target position estimation, and the simplified position track data base and tracking algorithm it allows, are planned to evaluate the trade-offs between timeliness and complexity of the two algorithms.

The limited sharing of information between nodes in the current test-bed does not interfere with position tracking by individual nodes. But it prevents nodes from "handing over" target position tracks as the targets pass through the DSN's sensor coverage and can cause each node to have a myopic view of the targets in the DSN's overall sensor coverage as illustrated in the previous section. Results of recent research^{9,10} into related problems suggest that it should be possible to form a complete surveillance picture in each data processor if each one transmits locally computed target state estimates to other data processors and if each uses the state estimates it receives as well as available raw sensor data to update its local target state estimates. But the research results demonstrated only asymptotic convergence of estimates of an invariant quantity. We must examine this work carefully to extract clues as to the proper organization of algorithms for our more complex situation. At the very least, we would like to develop an ad hoc algorithm for combining state estimates for the purposes of position track hand-over that exhibits minimal pathological behavior.

The need to demonstrate a system which performs satisfactorily in a realistic environment will continue to drive the development of the Lincoln Laboratory DSN test-bed tracking system. In parallel with the development of new tracking algorithms, data will be collected on actual sound sources in stressing environments such as crossing targets on known trajectories. Regular experimentation with such data focuses development of the system on significant problems, providing a starting point for theoretical investigations and a timely test of any ad hoc or approximate aspects of the system.

References

1. J. Capon, "High Resolution Frequency-Wavenumber Spectrum Analysis," Proc. of the IEEE, 57, No. 8, pp. 1408-1418 (August 1969).
2. "Distributed Sensor Networks Semiannual Technical Summary," Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA (30 September 1979) A086800.

3. "Distributed Sensor Networks Semiannual Technical Summary," Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA (30 September 1980) A103045.
4. R. P. Hughes, "A Distributed Multiobject Tracking Algorithm for Passive Sensor Networks," Master's Thesis, Massachusetts Institute of Technology (June 1980).
5. J. R. Delaney, "Location Algorithm Sensitivities," in preparation as Technical Note, Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA.
6. P. E. Green, "DSN Test-Bed Tour and Demonstration," Distributed Sensor Networks Workshop, Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA, pp. 163-174 (January 1982) A118868.
7. "Distributed Sensor Networks Semiannual Technical Summary," Lincoln Laboratory, Massachusetts Institute of Technology, Lexington, MA (31 March 1978) A060685.
8. R. L. Walton, "Multisite Acoustic Location," Distributed Sensor Networks Workshop, Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA, pp. 63-67 (December 1978).
9. V. Borkar and P. P. Varaiya, "Asymptotic Agreement in Distributed Estimation," IEEE Trans. on Automatic Control, AC-27, No. 3, pp. 650-655 (June 1982).
10. J. N. Tsitsiklis and M. Athans, "Convergence and Asymptotic Agreement in Distributed Decision Problems," Proc. of the 5th MIT/ONR Workshop on Distributed Information and Decision Systems, Monterey, CA (August 1982).

*This work was sponsored by the Defense Advanced Research Projects Agency.

The U.S. Government assumes no responsibility for the information presented.

Figures

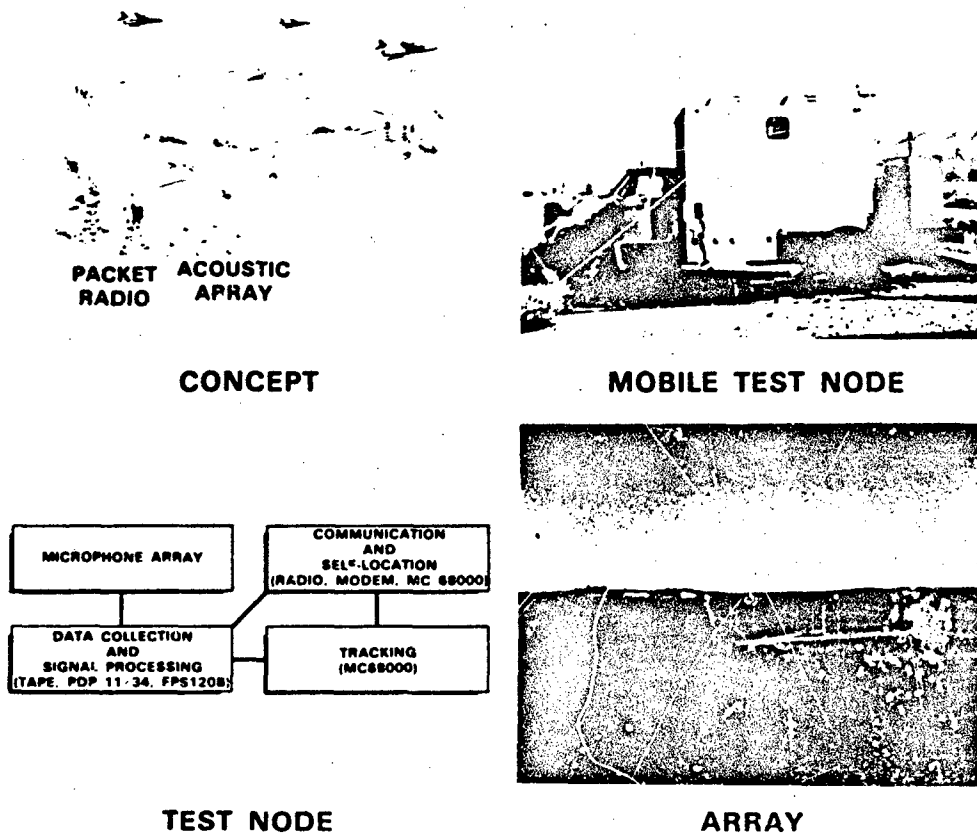


Fig. 1. The MIT/LL DSN Test-Bed.

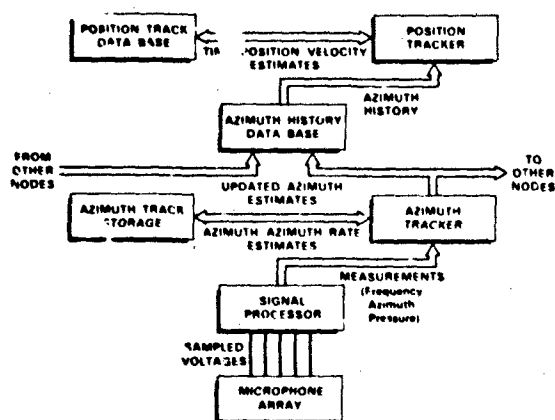


Fig. 2. Functional Description of the MIT/LL DSN Test-Bed Tracking System.

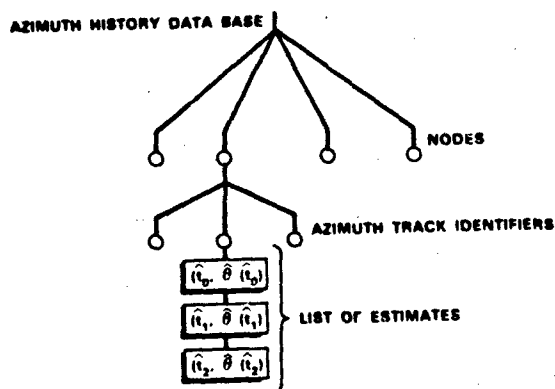


Fig. 3. Azimuth History Data Base Organization.

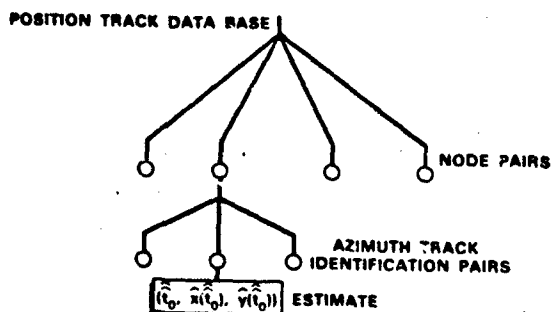


Fig. 4. Position Track Data Base Organization.




Fig. 5. UH-1 Flight Path for November 1981 Experiment.




Fig. 6. Azimuth Angle Measurements and Azimuth Tracks for Node F.




Fig. 7. Position Tracks Formed at Node F from Azimuth Histories Originating at Nodes F and H.




Fig. 8. Position Tracks Formed at Node J from Azimuth Histories Originating at Nodes H, J and L.

DISTRIBUTED ESTIMATION SYSTEMS*

C.Y. Chong, S. Mori

E. Tse

Advanced Information & Decision Systems
201 San Antonio Circle, Suite 286
Mountain View, CA 94040

Dept. of Engin.-Economic Systems
Stanford University
Stanford, CA 94305

ABSTRACT

In this paper, we consider the distributed estimation problem by a set of agents connected by an arbitrary communication network. The agents communicate conditional probabilities of the random state over the network. From these conditional probabilities, each agent then tries to re-construct the conditional probability given all the measurements if these were communicated instead of the probabilities. It is discovered that in general the agents have to remember some of the past conditional probabilities and may even have to request additional information. A method for generating the fusion algorithm for each agent based on the network structure is presented and applied to some examples. The results are applicable to both dynamic and static states.

1. INTRODUCTION

The traditional approach to estimation has been centralized. Even though the measurements are generated by a large number of sensors, it is usually assumed that they are sent to a central site where processing is carried out by one agent (computer). In this context centralized estimation theory is well developed and has found applications in many real world problems.

In recent years, there has been growing interest in distributed estimation problems. In such problems, the sensor measurements are not all transmitted to a central processor. Instead, a set of local processors, which we call estimation agents, are present. The agents are connected by a communication network. Each agent collects the measurements from a subset of the sensors, performs some local processing, and communicates the results with other agents.

The advantages of such a distributed estimation system are many. It is more reliable (or less vulnerable) since there is not a single central site which is responsible for the proper functioning of the system. Communication is cheaper since only the results of processing, and not the raw data, are communicated. Furthermore, each distributed agent has the use of the processed data locally and does not have to wait for communication from the central processor. From a technological point of view, such distributed systems are made possible by the availability of cheap computing hardware. These advantages make distributed estimation systems extremely attractive for many military and civilian applications. One such application is the distributed sensor network [1], [2] for tracking and surveillance.

*This work was supported by the Defense Advanced Research Projects Agency under Contract MDA-903-81-C-0333.

Research in distributed estimation has progressed along several directions. A team-theoretic approach has been taken by Barta [3] for decentralized linear estimation and by Tenney and Sandell [4] for distributed detection. Extensions of this work in detection have been made by Teneketzis [5] and Ekchian and Tenney [6]. Another approach, based on finding constrained decentralized filters, has been taken by Tacker and Sanders [7]. The approach of fusion or combining of local estimates to recover the globally optimal estimate has been used in [8] to [12]. The linear problem was considered by Speyer [8], Chong [9], Willsky et al. [10] and Levy et al. [11] while Castanon and Teneketzis [12] considered the nonlinear extension. In all of the above [8]-[11], the system structure is hierarchical with no feedback communication or coordination from the fusion agent. Similar problems of this type have also been considered in the management science literature [13].

The network aspect in the distributed estimation problem has been the emphasis in [14], [15] and discussed in [2]. Borkar and Varaiya [14] presented results on the asymptotic agreement among agents for estimation while Tsitsiklis and Athans [15] considered asymptotic agreement for more general decision problems. It has been demonstrated in [2] via an example that agreement may not be desirable since the common conclusion may be wrong.

In this paper, we elaborate the results obtained in [2]. The philosophy of fusion or combining of local conditional probabilities to obtain the probability conditioned on all available information is again used. However, arbitrary network structures are considered explicitly. They may be hierarchical with or without feedback from the higher level or fully distributed. The presentation is at a fairly elementary level to simplify the notation but can be made more sophisticated if desired by introducing sigma fields. The results may provide the theoretical basis for the analysis and design of systems such as the distributed sensor network.

The rest of this paper is organized as follows. In Section 2, we present the model to be used for distributed estimation. Section 3 describes the distributed estimation problem. Section 4 describes the basic results for static random states. A method for generating the fusion formula for arbitrary networks is given. The fusion algorithms for some examples are also described. Section 5 extends the basic results to the case of dynamic random states. Section 6 is the conclusion.

2. MODEL FOR DISTRIBUTED ESTIMATION

2.1 State and Observation Models

We consider the estimation of a random process $x(t)$, $t \in T$ where $T = [t_0, \infty)$ and $x(t) \in X$. The random process $x(\cdot)$ can be static, deterministic or a general

Markov process. We assume the statistics which specify the random process completely are known.

Let S be a finite set of sensors. At a given time t in T , a sensor s generates an output or measurement z in the measurement space Z_s . The triple (z, t, s) is then called a data set and $a(t, s)$ is the data set index. Let Z be the set of all data sets and K be the set of all data set indices. If we assume that each sensor can produce only a finite number of outputs in any finite time interval, the sets Z and K are at most countable. Furthermore, for each $t \in T$, the restrictions $Z_t = \{(z, t', s) | z | t' \leq t\}$ and $K_t = \{(t', s) | t' \leq t\}$ are both finite. We make two additional assumptions:

1. The sensor origin and time of each data set are known, i.e., for any data set $(z, t, s) \in Z$, t and s are known quantities.
2. The measurements are all conditionally independent given the state process, i.e., for any finite subset $\{(z_1, t_1, s_1), \dots, (z_k, t_k, s_k)\}$ of Z ,

$$\text{Prob.} \left(\bigcap_{i=1}^k \{z_i \in dz_i\} | x(t_1), \dots, x(t_k) \right) = \prod_{i=1}^k \text{Prob.}(z_i \in dz_i | x(t_i)) \quad (2.1)$$

With the second assumption, the observation process can be characterized completely by the transition probabilities (or probability densities) from X to Z_s .

2.2 Data Bases

We are interested in estimation of the process by a network of agents. At any time t , due to communication constraints, each agent may not have access to all available data sets. In general, an agent will have only a subset of the available Z_t at t , corresponding to only a subset of K_t . A data base Z at time t is a subset of Z_t and a data index base K at time t is a subset of K_t . According to this definition, Z_t (K_t) is the maximum data (index) base at t and ϕ (the empty set) is the minimum. Given any data base $Z = \{(z_1, t_1, s_1), \dots, (z_k, t_k, s_k)\}$, the corresponding data index base $K = \{(t_1, s_1), \dots, (t_k, s_k)\}$ is found by the operation $K = I(Z)$ where the definition of I_n is obvious and the actual measurements (z_1, \dots, z_k) are found by $(z_1, \dots, z_k) = M_n(Z)$. When $Z = \phi$, $I(\phi) = \phi$, and $M_n(\phi) = \theta$ where θ is a symbol representing "no information".

For each data index base $K = \{(t_1, s_1), \dots, (t_k, s_k)\}$ with corresponding data base $Z = \{(z_1, t_1, s_1), \dots, (z_k, t_k, s_k)\}$ we define the conditional probability $P(.|Z)$ to mean $P(.|M_n(Z), K)$.

All the definitions above can be given more rigorously in terms of sigma algebras. This will not be attempted in this paper so as to simplify the development.

2.3 Communication Model

We assume there is a finite set N of estimation agents. Each agent n has its own set of sensors, i.e., a subset S_n of S . Furthermore, the sensor sets are disjoint for different agents, i.e., $S_n \cap S_{n'} = \phi$ for $n \neq n'$. Each agent n also receives information from other agents via communication. Communication among agents is specified by the known communication schedule C which is a subset of $T \times N \times N$. $(t, n_1, n_2) \in C$ means that agent n_1 transmits some messages to agent n_2 at time t . The exact form of the messages will be discussed later.

3. DISTRIBUTED ESTIMATION PROBLEM

3.1 Information Graph

The distributed estimation system (N, S, C) thus consists of the sensor set S and the estimation agent set N together with the communication schedules. Four types of events affect the change of information in the system. These events, the times when they occur and the nodes (sensors or estimation agents) which are affected, are given below:

- sensor observation: K ,
- reception of sensor data by estimation agent: $\{(t, n) \in T \times N | (t, s) \in K, s \in S_n\}$,
- transmission by estimation agent: $\{(t, n) \in T \times N | (t, n, n') \in C\}$,
- reception of transmission by estimation agent: $\{(t, n) \in T \times N | (t, n', n) \in C\}$.

Consider a subset I of $T \times (S \cup N)$ which is the union of all the sets defined above. Define an anti-symmetric and transitive binary relation (or partial ordering) $<$ on I such that

- i. For each $(n, t, t') \in N \times T \times T$, $(t, n) \in I$, $(t', n) \in I$ and $t < t'$ implies that $(t, n) < (t', n)$;
- ii. $(t, s) \in K$, $s \in S_n$ and $(t, n) \in I$ implies that $(t, s) < (t, n)$;
- iii. $(t, n, n') \in C$ implies that $(t, n) < (t, n')$.

This binary relation or partial order on I thus satisfies all the constraints associated with perfect communication as defined by C as well as perfect memory at each processing node. $(I, <)$ characterizes the information flow in the system and is called the information graph. If all the sensor measurements (data sets) can be communicated perfectly through the communication network, the data base $Z(t, i)$ for each node (t, i) in the graph $(I, <)$ can be defined by beginning with the minimal elements and following the rules shown below:

- i. If (t, i) is a receiving node,
 $Z(t, i) = \{Z(s, j) | (s, j) \rightarrow (t, i)\}$;
- ii. If (t, i) is a transmitting node,
 $Z(t, i) = \begin{cases} Z(s, j) & \text{if } (s, j) \rightarrow (t, i) \\ \{(z(k), k)\} & \text{if } (t, i) = k \in K \\ \phi & \text{otherwise.} \end{cases}$

In the above $(s, j) \rightarrow (t, i)$ means that (s, j) is an immediate predecessor of (t, i) and $(z(k), k) \in Z$ is the unique element whose second component is k .

With this construction of the data base, we see that $(t, i) < (s, j)$ if and only if $Z(t, i) \subseteq Z(s, j)$. Similar remarks can be made for the data index base $K(t, i)$. Since there is a natural direction (along increasing time) in the graph, the arrowheads on the edges in a pictorial representation of the graph can be omitted. We would also omit those edges which are due to transitivity. From the graph, the flow of information in the system becomes very obvious. A node (t, i) is a parent of (s, j) if information flows from node i at time t to node j at time s . Note that in the information graph, the receiving nodes correspond to the events when estimates have to be updated with the arrival of new information. For many applications, it is sufficient to use a reduced information graph, which is obtained by considering only these receiving nodes.

Several examples of distributed estimation networks and their reduced information graphs are shown in Figures 1-4 where the hollow circles and the solid circles are the communication reception and sensor data

reception nodes respectively.

Example 1 (Fusion Without Coordination): Of the agents in N , agent 1 is a fusion agent and the rest are local agents. The local agents transmit to the fusion agent after they receive the data from the sensors and perform local processing. Figure 1 shows the structure of the system (for three agents) and the information graph.

Example 2 (Fusion With Coordination): This is similar to Example 1 except that right after fusion, agent 1 communicates with the local agents again. This structure is also equivalent to a broadcast system where all agents communicate with each other. Figure 2 shows the structure of the system and the information graph.

Example 3 (Cyclic Communication): This is the example considered in [2]. The agents are arranged in a circle so that each agent transmits only to its immediate neighbor in a cyclic manner at the specified communication times. Figure 3 shows the example for $N = \{1, 2, 3\}$.

Example 4 (Multipath Pattern): The agents are arranged as in Figure 4. agent 1 can only get information from agent 4 via agents 2 and 3.

3.2 Distributed Fusion Problem

The problem is to compute $p(x(t)|Z(t,i))$ for each node $(t,i) \in I$ in the graph $(I, <)$. Since the conditional probabilities or any estimates are updated only at the receiving nodes (extrapolation is carried out at the other nodes), we need only to consider the computations at the following two types of nodes in the reduced information graph: sensor data reception nodes and communication reception nodes.

At a sensor data reception node (t,i) , computation of $p(x(t)|Z(t,i))$ is straightforward. The standard Bayesian update formula would suffice. At a communication reception node, the objective is to reconstruct $p(x(t)|Z(t,i))$ from the conditional probabilities $\{p(x(t)|Z(s,j)) | (s,j) < (t,i)\}$. This problem is the distributed fusion problem: construction of the conditional probability given all the data sets which would have been communicated through the network using only the conditional probabilities available at the predecessor nodes in the information graph.

4. STATIC RESULTS

In this section we develop the main results for fusion for each agent i , assuming the random process is static, i.e., $x(t) = x$ for all t . Since the information from different agents may overlap, care has to be taken when the conditional probabilities from different agents are combined. In particular, any redundant information has to be identified so that it is not used more than once. The following lemmas provide the mechanism for doing this. In the following x denotes a random vector with prior probability $p(x)$ and Z is the set of all data sets.

4.1 Basic Results

We state the following lemmas without proofs, some of which can be found in [16].

Lemma 1: Suppose Z_1 and Z_2 are data bases at two information nodes 1 and 2. Then

$$p(x|Z_1 \cup Z_2) = C \frac{p(x|Z_1) p(x|Z_2)}{p(x|Z_1 \cap Z_2)} \quad (4.1)$$

where C is a normalization constant.

This lemma states that since $p(x|Z_1)$ and $p(x|Z_2)$ both include information contained in the data base $Z_1 \cap Z_2$, this common information has to be removed so that it does not get double counted. Lemma 1 plays a central role in distributed estimation theory similar to the usual Bayes' rule in centralized estimation theory. When the conditional probabilities from multiple agents are combined, the fusion formula can be obtained by repeated applications of Lemma 1. The following gives the results for three agents.

Lemma 2: Suppose Z_1, Z_2 and Z_3 are data bases at the information nodes 1, 2 and 3. Then

$$p(x|Z_1 \cup Z_2 \cup Z_3) = C \frac{p(x|Z_1 \cup Z_2) p(x|Z_3)}{p(x|(Z_1 \cup Z_2) \cap Z_3)} \\ = C \frac{p(x|Z_1) p(x|Z_2) p(x|Z_3) p(x|Z_1 \cap Z_2 \cap Z_3)}{p(x|Z_1 \cap Z_2) p(x|Z_2 \cap Z_3) p(x|Z_3 \cap Z_1)} \quad (4.2)$$

This lemma again has a very intuitive explanation. The terms in the denominator consist of pairwise redundant information to be removed. When these are removed, all information which is common to Z_1, Z_2 , and Z_3 is also removed. This then has to be restored.

If all the random elements involved are Gaussian, the lemmas above can be simplified so that only the conditional means and covariances are involved. Suppose x is Gaussian with mean m and covariance $P(0)$. Let $\hat{x}(Y)$ and $P(Y)$ be the mean and covariance corresponding to the conditional density $p(x|Y)$. Then lemma 1 becomes

Lemma 1A:

$$P(Z_1 \cup Z_2)^{-1} = P(Z_1)^{-1} + P(Z_2)^{-1} - P(Z_1 \cap Z_2)^{-1} \quad (4.3)$$

and

$$P(Z_1 \cup Z_2)^{-1} \hat{x}(Z_1 \cup Z_2) = P(Z_1)^{-1} \hat{x}(Z_1) + P(Z_2)^{-1} \hat{x}(Z_2) \\ - P(Z_1 \cap Z_2)^{-1} \hat{x}(Z_1 \cap Z_2). \quad (4.4)$$

Lemma 2 can be simplified in a similar way. Lemma 1A is identical to that used in [9] for deriving the optimal algorithms for combining estimates of linear Gaussian systems.

We now state the static fusion problem for each agent assuming that $x(t) = x$ for all t . The problem is stated for the case when messages are received from only one agent. But the extension to multiple agents is obvious.

Static Fusion Problem

Suppose agent i receives a message from agent j at time s in the form of a conditional probability $p(x|Z(s,j))$. Let (t,i) be the immediate predecessor to (s,i) for agent i . Agent i 's data base then changes from $Z(t,i)$ to $Z(s,i) = Z(t,i) \cup Z(r,j)$ where (r,j) is the immediate predecessor to (s,j) for agent j . The objective is to find $p(x|Z(s,i))$ in terms of $p(x|Z(t,i))$, $p(x|Z(r,j))$ and possibly other conditional probabilities defined on the information graph, i.e., $\{p(x|Z(t',i')) | (t',i') < (s,i)\}$.

We do not specify a priori which conditional probabilities are involved except that they have to be conditional on some data base Z defined on the information graph and that they should be available through communication. The following recursive algorithm allows us to find the set of needed conditional probabilities and

how they should be combined.

Algorithm for Static State

The algorithm consists of repeated applications of the following steps.

Step 1: Since $Z(t,i)$ and $Z(r,j)$ are subsets of Z , Lemma 1 gives

$$p(x|Z(s,i)) = p(x|Z(t,i) \cup Z(r,j)) = c \frac{p(x|Z(t,i)) p(x|Z(r,j))}{p(x|Z(t,i) \cap Z(r,j))} \quad (4.5)$$

If $Z(t,i) \cap Z(r,j)$ is the data base for some node in the information graph, i.e., $Z(t,i) \cap Z(r,j) = Z(q,k)$ for some (q,k) in I or if it is empty, then the algorithm terminates. If not, Step 2 is used. In terms of the information graph representation introduced in Section 3, this step is particularly simple. We start from two information nodes (t,i) and (r,j) whose conditional probabilities are to be combined. $Z(t,i) \cap Z(r,j)$ corresponds to the information of all those nodes which are parents of both (t,i) and (r,j) .

Step 2: Let $\{(t_1,k_1), (t_2,k_2), \dots\}$ be the set of common predecessors of (t,i) and (r,j) in the information graph. Then

$$Z(t,i) \cap Z(r,j) = Z(t_1,k_1) \cup Z(t_2,k_2) \cup \dots \quad (4.6)$$

Step 1 can now be repeated with the help of Lemma 1 (and its multiple agent version) to express $p(x|Z(t,i) \cup Z(r,j))$ in terms of the conditional probabilities $p(x|Z(t,k_1))$, $i = 1, 2, \dots$, and $p(x|Z(t,k_j) \cap Z(t,k_1))$, $i = 1, 2, \dots$, $j = 1, 2, \dots$, etc. The algorithm terminates when all the conditional probabilities are defined on nodes in the information graph or coincide with the a priori distributions.

By applying this algorithm, $p(x|Z(t,i) \cup Z(r,j))$ can be expressed in terms of products and ratios of conditional probabilities defined on information nodes. Each product corresponds to the fusion or combining of information whereas each division corresponds to the removal of redundant information. Note that in general it is not sufficient to use only the conditional probabilities $p(x|Z(t,i))$ and $p(x|Z(r,j))$ unless $Z(t,i)$ and $Z(r,j)$ happen to be disjoint or there is a node (s,k) such that $Z(s,k) = Z(t,i) \cap Z(r,j)$. Additional conditional probabilities from the past are also needed so that the redundant information in $Z(t,i)$ and $Z(r,j)$ can be identified and removed.

We have thus solved the fusion problem for each agent in a distributed estimation network. This algorithm also provides us with the set of conditional probabilities which needs to be stored at each agent plus the additional set of conditional probabilities which needs to be communicated.

When the random elements involved are all Gaussian, the sufficient statistics for the conditional probabilities become the conditional means and covariances. With the help of Lemma 1A, we can again apply the algorithm. Instead of multiplication and division of probabilities, however, we now have operations involving conditional means and covariances. The results are straightforward and will not be presented here.

4.2 Static Examples

In the following we assume the measurements are made at times $\{\dots, t-1, t+1, \dots\}$ and messages are received at times $\{\dots, s-1, s+1, \dots\}$ with $s-1 < t < s$.

Example 1 (Fusion Without Coordination): Consider the fusion time s . Let t be the observation time immediately before s . With the information graph it is easy to see that $Z(s-1,1) \cap Z(t,2) = Z(t-1,2)$. Thus

$$p(x|Z(s-1,1) \cup Z(t,2)) = c \frac{p(x|Z(t,2)) p(x|Z(s-1,1))}{p(x|Z(t-1,2))}$$

By a recursive argument, we can show that

$$p(x|Z(s,1)) = c \prod_{i=1}^s \frac{p(x|Z(t,i))}{p(x|Z(t-1,i))} p(x|Z(s-1,1)). \quad (4.7)$$

Each term in the product contains the new information contained in the new measurement $z(t,i)$ of agent i . All other information is already known to agent 1. The fusion problems of the other agents are similar.

Example 2 (Fusion with Coordination): This is equivalent to broadcast communication. From the information graph, the algorithm gives for j

$$p(x|Z(s,j)) = c \prod_i \frac{p(x|Z(t,i))}{p(x|Z(s-1,i))} p(x|Z(s-1,j)). \quad (4.8)$$

Each term in the product is the new information contained in measurement $z(t,i)$.

Example 3 (Cyclic Communication): The algorithm gives for general $i = 1, 2, 3$

$$p(x|Z(s,i)) = c \frac{p(x|Z(t,i))}{p(x|Z(t-2,i))} \frac{p(x|Z(t,[i+1]))}{p(x|Z(t-1,[i+1]))} p(x|Z(s-3,i)) \quad (4.9)$$

where $[i]$ is i modulo 3.

Thus, in addition to the most current conditional probability $p(x|Z(t,1))$, agent 1 has to remember three other probabilities. Note that $p(x|Z(t-1,2))$ is available to agent 1 from earlier communications. This indicates that in a distributed estimation network, knowing the most recent estimate is frequently not sufficient if one wants to recover the globally optimal estimate. In fact, it has been shown via simulation in [2] that if a suboptimal rule of combining estimates is used, such as

$$p(x|Z(t,1) \cup Z(t,2)) \approx c p(x|Z(t,1)) p(x|Z(t,2)) \quad (4.10)$$

for agent 1 and similar rules for agents 2 and 3, the agents agree asymptotically. This is consistent with the results on asymptotic agreement in distributed estimation as given in [14]. However, the agents can converge to the wrong estimate as demonstrated in [2]. Thus, although optimal fusion algorithms are in general more complicated, requiring more memory and more computation, they are nonetheless necessary if good performance is needed. A suboptimal algorithm has also been tested in [2] and shown to have some nice properties.

Example 4 (Multipath Pattern): The fusion problems of agents 2 and 3 are straightforward. For agent 1, repeated use of the algorithm (with the help of the information graph in Figure 4) gives

$$p(x|Z(s,1)) = c \frac{p(x|Z(t,2))}{p(x|Z(t-1,2))} \frac{p(x|Z(t,3))}{p(x|Z(t-1,3))} \frac{p(x|Z(t-2,4))}{p(x|Z(t-1,4))} p(x|Z(t,1)) \quad (4.11)$$

In addition to the conditional probabilities from agents 2 and 3, conditional probabilities by agent 4 are also needed. These would have to be relayed by agents 2 or 3.

In the above examples, general fusion formulas are given. If the random vectors are all Gaussian, these formulas can be simplified using Lemma 1A.

5. DYNAMIC RESULTS

Assume now that $x(\cdot)$ is a Markov process. The fusion problem for each agent will now be considered. Since the data sets are no longer conditionally independent given $x(t)$, one immediate question is the choice of an appropriate "state" whose conditional probabilities would be computed, transmitted and combined by the various agents. Let $T(t, i)$ be

$$T(t, i) = \{t' \in T(t', i') \in K(t, i)\}, \quad (5.1)$$

and

$$y = (x(t'))_{t' \in T(t, i)} \quad (5.2)$$

for each information node (t, i) where fusion is to be performed. Then the problem is effectively reduced to a static problem of the type considered in Section 4. Using the independence assumptions on the measurements in the data base given y , the algorithm in Section 4 can be applied. However, this means that the conditional probability of a high dimensional random vector y would have to be stored and transmitted. From an implementational point of view, this may not be feasible.

For deterministic random processes, which can be characterized by the state at one given time, an obvious choice is to estimate $x(t_0)$ where t_0 is the minimum in the set T . Again, due to the Markov property, the conditional independence assumption is satisfied and the algorithm can be used. However, if there are substantial changes in the process, $x(t_0)$ may not be the state of interest. In this section, we characterize the more current states whose conditional probabilities ought to be transmitted and combined.

The following generalization of Lemma 1 is needed.

Lemma 3: Consider a random vector y and data bases Z_1 and Z_2 defined on the information graph. Suppose

$$\begin{aligned} p(Z_1 - Z_2, Z_2 - Z_1 | Z_1 \cap Z_2, y) \\ = p(Z_1 - Z_2 | y, Z_1 \cap Z_2) p(Z_2 - Z_1 | y, Z_1 \cap Z_2). \end{aligned} \quad (5.3)$$

Then

$$p(y | Z_1 \cup Z_2) = C \frac{p(y | Z_1) p(y | Z_2)}{p(y | Z_1 \cap Z_2)} \quad (5.4)$$

where C is a normalization constant and $A-B$ denotes the difference of the sets A and B .

Lemma 3 states that even though the individual measurements in Z do not satisfy the conditional independent assumptions given y , Equation (5.4) (which is the same as (4.1)) is still valid provided the private data bases $Z_1 - Z_2$, $Z_2 - Z_1$ are conditionally independent given the state y and the common information $Z_1 \cap Z_2$.

We can now state the following theorem which characterizes the state vector which should be estimated for deterministic dynamic random processes.

Theorem: Consider the fusion problem for the information node (t, i) assuming a deterministic random process x . If the algorithm of Section 4 yields the fusion formula

$$p(x | Z(t, i)) = F(p(x | Z(t', i'))); (t', i') \in L(t, i) \quad (5.5)$$

where F is a function consisting of products and ratios of $p(x | Z(t', i'))$'s in the set $L(t, i)$, and $L(t, i)$ is a subset of the predecessor information nodes of (t, i) .

Then for a deterministic random process $x(\cdot)$, equation (5.5) holds with x replaced by $x(t^*)$, where

$$t^* = \min\{t' | (t', s) \in L(t, i) - \{(t', s^*)\}\} \quad (5.6)$$

and (t', s^*) is the minimal element in $L(t, i)$.

The proof is straightforward and is based on the algorithm of Section 4 and Lemma 3. This theorem states that for random processes, in general the filtered estimate represented by the conditional probabilities $p(x(t) | Z(t, i))$ may not be adequate for optimal fusion at time t . Sometimes the agents need to have the conditional probabilities of the states at some earlier times. Thus, smoothed estimates are frequently needed. From this, the estimates of the current states can be obtained easily by extrapolation. When this theorem is applied to the examples in Section 4, we obtain the following results.

Example 1 (Fusion without Coordination): In the fusion equation (4.7), the state to be estimated is $x(t)$. This is consistent with the results in [8]-[12]. As a variation of this, consider a periodic fusion situation where the local agents acquire measurements at a higher rate than they communicate with the fusion agent (Figure 5). Specifically, let the new fusion time set for agent 1 be $\{\dots, s-M, s, s+M, \dots\}$ where M is the number of time units between communication. Then application of the theorem yields

$$\begin{aligned} p(x(t-M+1) | Z(s, 1)) = C \prod_{i=1} p(x(t-M+1) | Z(t, i)) \\ p(x(t-M+1) | Z(s-M, 1)) \end{aligned} \quad (5.7)$$

Thus, the state of interest is now $x(t-M+1)$, and each term in the product contains the new information of agent i about this state.

Example 2 (Fusion with Coordination): In equation (4.8), the state is $x(t)$.

Example 3 (Cyclic Communication): In equation (4.9), the state is $x(t-2)$. Thus, extrapolation is needed if the estimate of $x(t)$ is needed.

Example 4 (Multipath Pattern): In equation (4.11), the state is $x(t-1)$. Thus, extrapolation is again needed.

6. CONCLUSION

We have presented a formalism for the distributed estimation problem. Using this formalism, the optimal fusion algorithm for each agent in the network has been developed for arbitrary network structures. Both results for static and deterministic dynamic random states have been described, and illustrated with examples. The results have been presented for very general state and observation models. Special cases such as linear models with Gaussian noises can be considered. An interesting special case for distributed multitarget tracking and classification has also been investigated and briefly reported in [2]. The details will appear elsewhere.

REFERENCES

1. R. Iacoss, "Overview of the Distributed Sensor Networks Program at Lincoln Laboratory," *Proc. Distributed Sensor Networks Workshop*, MIT Lincoln Lab. Jan. 6-7, 1982.
2. C.Y. Chong, S. Mori, E. Tse and R.P. Wishner, "Distributed Estimation in Distributed Sensor Networks," *Proc. 1982 American Control Conference*, Arlington, VA, 1982.

3. S.M. Barta, "On Linear Control of Decentralized Stochastic Systems," Ph.D. Dissertation, MIT, Cambridge, MA, July 1978.
4. R.R. Tenney and N.R. Sandell, Jr., "Detection with Distributed Sensors," *IEEE Trans. Aerospace and Electronic Systems*, Vol. AES-17, No. 4, pp. 501-510, July 1981.
5. D. Teneketzis, "The Decentralized Quickest Detection Problem," *Proc. 21st IEEE Conf. on Decision and Control*, Fort Lauderdale, FL, 1982.
6. L.K. Ekchian and R.R. Tenney, "Detection Networks," *Proc. 21st IEEE Conf. on Decision and Control*, Ft. Lauderdale, FL, 1982.
7. E.C. Tacker and C.W. Sanders, "Decentralized Structures for State Estimation in Large Scale Systems," *Large Scale Systems Theory and Applications*, Vol. 1, No. 1, pp. 39-49, Feb. 1980.
8. J.L. Speyer, "Computation and Transmission Requirements for a Decentralized Linear-Quadratic-Gaussian Control Problem," *IEEE Trans. Automat. Contr.*, Vol. AC-24, No. 2, pp. 266-269, April 1979.
9. C.Y. Chong, "Hierarchical Estimation," *Proc. 2nd MIT/ONR Workshop on Distributed Communication and Decision Problems*, July 1979.
10. A.S. Willsky, M. Bello, D.A. Castanon, B.C. Levy and G. Verghese, "Combining and Updating of Local Estimates and Regional Maps Along Sets of One-Dimensional Tracks," *IEEE Trans. Automat. Contr.*, Vol. AC-27, No. 4, pp. 799-813, Aug. 1982.
11. B.C. Levy, D.A. Castanon, G.C. Verghese and A.S. Willsky, "A Scattering Framework for Decentralized Estimation Problems," *Automatica*, vol. 19, no. 4, pp. 373-394, July, 1983.
12. D.A. Castanon and D. Teneketzis, "Nonlinear Data Fusion," *Proc. 5th MIT/ONR Workshop on C Systems*, Monterey, CA, 1982.
13. R.L. Winkler, "Combining Probability Distributions from Dependent Information Sources," *Management Science*, Vol. 27, No. 4, pp. 479-486, April 1981.
14. V. Borkar and P. Varaiya, "Asymptotic Agreement in Distributed Estimation," *IEEE Trans. Automat. Contr.*, Vol. AC-27, No. 3, pp. 650-655, June 1982.
15. J.N. Tsitsiklis and M. Athans, "Convergence and Asymptotic Agreement in Distributed Decision Problems," *Proc. 21st IEEE Conf. on Decision and Control*, Ft. Lauderdale, FL, 1982.
16. C.Y. Chong, E. Tse, and S. Mori, "Distributed Estimation in Networks," *Proc. 1983 American Control Conf.*, San Francisco, CA, 1983.




Figure 1 Fusion Without Coordination




Figure 2 Fusion With Coordination




Figure 3 Cyclic Communication




Figure 4 Multipath Pattern




Figure 5 Periodic Fusion Without Coordination

OPTIMAL MANEUVER DETECTION AND ESTIMATION IN MULTIOBJECT TRACKING*

Dr. Thomas Kurien

Dr. Alan Blitz

Dr. Robert Washburn

ALPHATECH, Inc.
2 Burlington Executive Center
111 Middlesex Turnpike
Burlington, Massachusetts 01803

Dr. Alan Willsky

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

1. Introduction

The problem of multiobject tracking has been studied extensively over the past few years. Good summaries of these approaches are available in the survey papers by Reid [1] and Bar-Shalom [2]. There are two drawbacks that are common to all the algorithms that have been developed and studied so far. Firstly, the relationship of the algorithm to the optimal (Bayesian) algorithm cannot be clearly seen due to the ad hoc and sometimes arbitrary approximations that are introduced into the algorithm. Secondly, none of the algorithms explicitly model or account for maneuvering targets. Several algorithms [3]-[7] have studied maneuver detection and estimation for single targets. However, the contradictory requirements imposed by maneuvers and clutter for the selection of gate size has prevented any of these suboptimal algorithms from being extended to the multiobject tracking problem.

In this paper we develop an algorithm that overcomes both these drawbacks. We formulate the multiobject tracking problem within the framework of a hybrid state estimation problem. This permits the construction of the optimal solution to the multiobject tracking problem. However, due to the exponentially growing storage and computational needs of the algorithm with time, some form of suboptimal approximations have to be made. Since these approximations are made within the framework of the optimal algorithm, the nature and trade-offs associated with the approximations can be examined. Furthermore, optimal and suboptimal algorithms for maneuver detection and estimation can also be incorporated within this framework.

The paper is organized as follows. In Section 2 we formulate the problem of multiobject tracking of maneuvering targets within the framework of the hybrid state estimation problem. The computer implementation of the optimal solution is discussed in Section 3. Suboptimal features that reduce the computational requirements are discussed in Section 4. Finally, in Section 5, we provide some simulation results.

2. Problem Formulation and Optimal Solution

A multiobject tracking algorithm involves two basic functions - association of measurements with

postulated targets and utilization of the measurements to track the targets. The association of the measurements with postulated targets can be viewed as a discrete-valued state estimation problem. Subsequent to making these associations, tracking the targets corresponds to a continuous-valued state estimation problem. As

The model for a hybrid system can be represented as

$$\underline{x}(k+1) = A(k, q(k)) \underline{x}(k) + \underline{\zeta}(k, q(k)) \quad (1)$$

$$\underline{z}(k) = C(k, q(k)) \underline{x}(k) + \underline{\eta}(k, q(k)) \quad (2)$$

where

\underline{x} represents an n -dimensional continuous-valued state vector.

\underline{z} represents an m -dimensional measurement vector.

A is an $n \times n$ matrix representing the transition dynamics.

C is an $m \times n$ matrix representing the measurement process.

$\underline{\zeta}$ is an n -dimensional white-noise process representing model uncertainty.

$\underline{\eta}$ is an m -dimensional white-noise process representing measurement uncertainty.

q represents a discrete-valued stochastic process which take on l values.

k represents the time index.

Notice that the matrices A and C and the characteristics of the noise process are controlled by the process q . Their dimensions n and m are, in general, dependent on time. The discrete-valued process is assumed to be a Markov process with an l which is also time-dependent.

As mentioned earlier, the key features of a multiobject tracking algorithm in a track-while-scan system can be captured by this model. For example, \underline{x} represents the composite state of all postulated targets and \underline{z} represents the composite vector of all measurements at scan k . Further, the value of $q(k)$ specifies the

*This work was supported by the Department of the Navy, Naval Air Systems Command, under contract number N00019-C-82-0456.

association of measurements with postulated models through the matrices A and C. We will elaborate this in Section 3.

An optimal estimate of the state of the hybrid system, represented by Eq. 1 using the measurements in Eq. 2, can be obtained by computing the joint posterior probability density

$$p(\underline{x}(k), q^k | \underline{z}^k) \quad (3)$$

where

$$\underline{z}^k = (\underline{z}(1), \underline{z}(2), \dots, \underline{z}(k))$$

and

$$q^k = (q(1), q(2), \dots, q(k))$$

For most hybrid systems of practical importance, the transitions of the discrete-valued state are independent of $\underline{x}(k)$.^{*} In such cases the probability of a particular sequence q^k (which we will refer to as a hypothesis) can be computed recursively as

$$p(q^k | \underline{z}^k) = \frac{1}{C_k} [p(\underline{z}(k) | q^k, \underline{z}^{k-1}) p(q(k) | q^{k-1}, \underline{z}^{k-1}) p(q^{k-1} | \underline{z}^{k-1})] \quad (4)$$

where

C_k represents the summation of the numerator over all possible q^k .

The marginal density for $\underline{x}(k)$ can then be computed recursively from

$$p(\underline{x}(k) | \underline{z}^k) = \sum_{q^k} p(\underline{x}(k) | q^k, \underline{z}^k) p(q^k | \underline{z}^k) \quad (5)$$

where

$$p(\underline{x}(k) | q^k, \underline{z}^k) = \frac{1}{C_k} p(\underline{z}(k) | q^k, \underline{x}(k), \underline{z}^{k-1}) p(\underline{x}(k) | q^k, \underline{z}^{k-1}) \quad (6)$$

To summarize, the solution of the hybrid-state estimation problem involves the computation of the conditional density functions in Eq. 6 along with the probabilities of the hypotheses using Eq. 4. This enables the computation of the posterior probability of $\underline{x}(k)$ given in Eq. 5.

For the multiobject tracking problem at hand, if we assume that the dynamics of the individual targets (and measurement process) are linear and the noise processes are Gaussian, then the conditional density

^{*}Such hybrid systems have also been referred to as Event-Driven Dynamic systems or Dynamic Systems in a Switching Environment [8].

functions in Eq. 6 can be obtained by computing their sufficient statistics using a bank of Kalman filters. Now the hypothesis q^k represents a particular sequence of target and measurement associations. If a systematic method is used to compute these probabilities in Eq. 4, then the sufficient statistics for constructing the posterior probability of $\underline{x}(k)$ will be available. Using this, either a Maximum-a-Posteriori (MAP) or a Minimum-Variance-of-Error estimator can be constructed. In the next section we discuss how the probabilities in Eq. 4 can be reconstructed for the multiobject tracking problem.

3. Implementation of the Optimal Algorithm

As pointed out earlier, the essential difficulty in developing the optimal algorithm lies in constructing all the possible combinations of targets with measurements and then computing their probabilities. The combinatorial problem is even more aggravating because of changes in

1. number of targets due to births and deaths,
2. dynamic models of targets due to maneuvers, and
3. measurement characteristics due to clutter or missed measurements.

Two approaches recommended in the past [9] have attempted to construct the hypotheses in the form of a matrix. In one of these approaches, referred to as the target oriented approach [1], the postulated targets define the columns of the matrix and the postulated hypotheses define the rows. The entries of the matrix represent measurements. Then for a given row (hypothesis), the column numbers and the measurements in the associated columns specify the target-measurement pairs postulated by that hypothesis. A typical hypotheses matrix is shown in Fig. 1. The '0' entries indicate that the target is not detected.

		TARGET NUMBER		
		1	2	3
HYPOTHESIS NUMBER	1	1	2	3
	2	1	2	0
	3	1	0	3
	4	1	0	0
	5	0	2	3
	6	0	2	0
	7	0	0	3
	8	0	0	0

Figure 1. Hypotheses Matrix for Target-Oriented Approach

In the alternate approach, referred to as the 'measurement oriented approach,' the roles of the targets and measurements are interchanged. A typical hypotheses matrix using this approach is shown in Fig. 2. Here the '0' entries denote that the measurements corresponding to those columns are assumed to be false alarms.

		MEASUREMENT NUMBER		
		1	2	3
HYPOTHESIS NUMBER	1	1	2	3
	2	1	2	0
	3	1	0	3
	4	1	0	0
	5	0	2	3
	6	0	2	0
	7	0	0	3
	8	0	0	0

Figure 2. Hypotheses Matrix for Measurement-Oriented Approach

Both approaches have drawbacks. For example, in the target-oriented approach measurements, not included in a row, could correspond to either new targets or false alarms; this cannot be shown explicitly. Similarly, in the measurement-oriented approach, targets not included in a row could either have died or were not detected. The hypothesis matrix cannot display it. Furthermore, neither of the approaches can account for target maneuvers.

To overcome these problems, we have chosen to create the hypotheses at any scan in a novel fashion which is also intuitively appealing. Rather than representing the hypotheses in the form of a matrix, this approach maintains a set of target trees and a list of global hypotheses. The root of each target tree represents the birth of the target and the branches represent the different dynamics that the target can assume and the various measurements it can be associated with in subsequent scans. A trace of successive branches from a leaf to the root of the tree corresponds to a potential track of the target. The leaf of each such trace is unique and is referred to as a track node of the target tree.


Each element of the global hypotheses list contains a set of pointers which point to track nodes. They represent the combination of track nodes postulated by the global hypothesis which that element represents. Obviously, the collection of pointers in any one such global hypothesis cannot point to two track nodes within the same target tree.

The creation of the global hypotheses using target trees and global hypotheses list, enables the decomposition of the process of associating targets with measurements into that of associating measurements with each of the targets and then forming combinations of the resulting tracks. As such, we refer to this as a Track-Oriented approach. The expansion of the individual tracks at any scan can, in turn, be done in two stages. First, the tracks are split for possible dynamics and next these tracks are associated with the measurements. By assuming that the target dynamics are independent of the measurement characteristics, the transition diagrams for each of the targets and the measurements will have the simple form described below.

The discrete states and the associated transition diagram for a single maneuvering target is considered first. The target starts off in an unborn state (B),

is born at some scan and can then die (X) at some later scan. A target that is in the born state can have a constant velocity (nonmaneuver state S) or h_2 accelerating (maneuver state M). To allow for different accelerations that the target can undergo, there could be several maneuver states M_i ($i=1, n_m$). This is depicted in Fig. 3 where we have considered the case where $n_m=2$. For convenience in representing the transition diagram, we have made the following assumptions:


1. The target is always born into the non-maneuver state.
2. The probability of transitioning to any of the born states or X is independent of the current born state of the target.



R-1482

Figure 3. Transition Diagram for Target Dynamics States

The transition diagram associated with the measurement process is shown in Fig. 4. Observe that the probability of transitioning to either state is independent of the prior state. To prevent the existence of targets that have never been detected, we assume that a target that is born in the current scan will be detected. An alternate way of defining this requirement is to define the number of births parameter (in the distribution assumed for births) conditioned on the event that it will be detected. This also implies that the number of births conditioned on the event that it will not be detected is assumed to be zero.



R-1272

Figure 4. Transition Diagram for Measurement States

Now we can depict the construction of the global hypotheses in any scan. As mentioned, the track nodes of all target trees are extended in two steps.

In step 1 the track node is split into several branches - to account for each of the several dynamic models that the target can assume. This is shown in Fig. 5. Obviously, a parent track node that corresponds to a dead track is not split; only a continuation of the dead status is shown in this case.

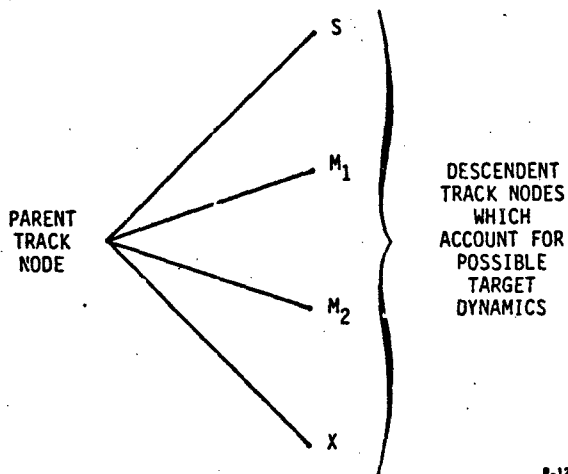


Figure 5. Track Splitting to Account for Different Target Dynamics

In step 2, the extended track nodes (excluding those which correspond to dead tracks) are associated with the measurements received in that scan. New track nodes are also generated to account for the possibility of a missed detection. Hence, if there are n_r returns in the scan, then each of the track nodes will have $(1 + n_r)$ descendants. We have extended the tree in Fig. 5 to illustrate the effect that step 2 has on the track splitting process for the case where $n_r = 2$ (Fig. 6). It is easy to see that for the general case, the number of track descendants for a maneuvering target is

$$[1 + (1+n_m)(1+n_r)] \quad (7)$$




Figure 6. Tracking Splitting to Account for Different Target Dynamics and Different Measurement Associations

Now we form all combinations of track nodes, which are descendants of parent track nodes included in the parent global hypothesis list, with the restriction that no two track nodes included in a new global hypothesis list should have the same parent node or use the same measurement from the current scan. We can show that for a parent global hypothesis which postulates the existence of n_t tracks, the number of descendant global hypotheses is

$$\sum_{i=0}^{\min(n_t, n_r)} \left(1 + \frac{n_m}{2}\right)^{n_t-i} \left(1 + n_m\right)^i 2^{(n_r + n_t - 2i)} \cdot \frac{n_r! n_t!}{(n_t-i)! (n_r-i)! i!} \quad (8)$$

where as defined before

n_m = number of possible maneuvers

n_r = number of returns in current scan

Likelihood Computation:

The likelihood of any descendant global hypothesis has been shown to be

$$P[q^k | z^k] = \frac{1}{C_k} [P(q^k | q^{k-1}, z^{k-1}) p(z^k | q^k, z^{k-1}) P(q^{k-1} | z^{k-1})] \quad (9)$$

where the subscript i denotes the specific hypothesis. Since the likelihoods are used as a basis for comparing the various global hypotheses, we can ignore the denominator - it being the same for all. The first term in the numerator represents probabilities of transitioning from the parent global hypothesis to each of the descendant global hypotheses. Posterior likelihoods of these tracks, after associating them with the different measurements available in the scan, are represented by the second term. Finally, the last term is the likelihood of the parent global hypothesis.

If the likelihood of a false alarm is normalized to unity, the remaining measurement association likelihoods can be scaled accordingly. In such a case, we need only consider the likelihoods for the track nodes shown in Fig. 6 for each of the targets. This makes it possible to compute the likelihood of a descendant global hypothesis following the same steps used for constructing it.

The state transition diagram for the target dynamics (Fig. 5) defines the transition probabilities between different target states. The posterior likelihoods of the measurement associations can be obtained from a Kalman filter after being premultiplied by the probability of detection P_D . The tracks which are postulated as being missed are multiplied by $(1-P_D)$ only. Thus, the likelihoods of each of the descendant track nodes can be computed. Then, after the proper descendant track nodes have been selected, the likelihood of the descendant global hypothesis can be computed as a product of the likelihood of the parent global hypothesis and the likelihoods of all the descendant track nodes included in it.

4. Suboptimal Techniques

The main purpose for designing suboptimal techniques is to reduce the computational burden associated with the optimal algorithm. Within the context of the optimal algorithm that we have constructed above, the computational burden can be reduced by either discarding some of the unlikely global hypotheses or using some computationally simpler algorithm for estimating the continuous-valued states. We will discuss only the former; several standard suboptimal techniques for the latter can be found in the literature (e.g., α - β tracker, constant gain Kalman filter).

Techniques available for reducing the number of global hypotheses can be grouped into one of the following

1. Screening
2. Pruning
3. Merging
4. Clustering

Both screening and pruning use the likelihoods to determine whether hypotheses may be discarded. Merging corresponds to the process of combining similar hypotheses. Grouping hypotheses in order to process them independently is referred to as Clustering. Since the optimal algorithm, described in Section 3, constructs the global hypotheses in two stages, these hypotheses reducing techniques can be applied during either the track expansion process or the global hypotheses building stage.

Screening techniques prevent less likely hypotheses from being formed or discard them after they are partially formed. We have incorporated several such options in the optimal algorithm. The first one is that of creating gates around track nodes and testing whether a measurement falls within this gate prior to forming a new descendent track node. Since the gate sizes tend to be large at the time of track initiation, an additional screening option has been provided. This is to not initiate maneuvers in target dynamics until its track is "well established." By well established tracks we mean tracks for which the velocity uncertainty is below a certain threshold. This screening technique will prevent the inclusion of tracks which postulate maneuvers as a consequence of the large gate sizes at the time of track initiation. If a target were to maneuver at the time of birth, it will be picked up as a new target with little loss in information caused by dropping its previous track.

Two other screening options that we have introduced are based on the physical limitations of the target. One takes into account the finite velocities that a target can have; the initial uncertainty of velocity states has been chosen to reflect this. The other option takes into account the finite accelerations that are feasible for a target; we have restricted the target from executing several different maneuvers in succession. In terms of the transition diagram shown in Fig. 3, this restriction implies that once a target enters a maneuver state, it can either remain in that state or return to the constant velocity state.

Pruning techniques discard hypotheses after they are formed. It can be affected in two ways: either deleting hypotheses which have a likelihood below a certain threshold or by limiting the global hypotheses at any stage to a fixed number. The former is difficult to design since the threshold will, in general, be

time varying. The latter technique is simpler to implement since there are no thresholds to be designed. We have incorporated the second option into our algorithm.

5. Simulation Results

Due to the large computational requirements of the optimal algorithm, it is not feasible to run any test scenario for more than 2 or 3 scans. Hence, we have run the algorithm with both screening and pruning options, discussed in Section 4, in effect. Two test cases that were simulated are described below.

Test Case 1

We have considered the simple case of a single target having the trajectory shown (indicated by the continuous line) in Fig. 7. The target starts with a heading of 30° . At scan 5, it executes a maneuver (-30° turn) and thereafter maintains a heading of 0° . We have generated clutter at every scan represented by \square s. Figure 7 indicates the location of clutter at each scan.




Figure 7. Test Case 1, Rank 1 Global Hypothesis

We have summarized the simulation parameters in Table 1 and the tracking algorithm parameters in Table 2. The heuristics that have been used to reduce the computational requirements are given in Table 3. Using the measurement noise specified in Table 1, it can be shown that the uncertainty in the velocity estimates will be reduced to less than 10 m/sec after 5 scans. This ensures that the heuristic that initiates maneuver hypotheses only after tracks are well established, will postulate maneuvers for the target prior to the actual maneuver at scan 5.

Figure 7 also shows the trajectories postulated by the global hypothesis with the highest likelihood (highest rank). It can be seen that it identifies the correct target trajectory (dotted line through the Δ s). However, it postulates the existence of another target (dotted line through the \square s) that is born, detected, not detected and dead in successive scans starting with the fourth scan. This is a consequence of the large gates associated with the targets that are just born. At scan 9 such a track is insignificant and hence can be ignored.

TABLE 1. SIMULATION PARAMETERS

Scan Time T: 10 secs.

Number of Scans: 9

Surveillance Area: $-15000 < x < 15000$

$-15000 < y < 15000$

Measurement Noise: $\sigma_x = 30$ m

$\sigma_y = 300$ m

Initial Velocity of Target: (Speed: 300 m/s, 30° heading)

Target Velocity after 5th Scan: (Speed: 300 m/s, 0° heading)

Clutter: 1 per scan

Uniform between $\begin{bmatrix} (x_t(k) - 1500) \text{ and } (x_t(k) + 1500) \\ (y_t(k) - 15000) \text{ and } (y_t(k) + 15000) \end{bmatrix}$

TABLE 2. PARAMETERS USED IN TRACKING ALGORITHM

Initial Filter Covariance: $\text{Diag} [p_{11}, p_{22}, p_{33}, p_{44}]$

p_{11}, p_{22} : Set based on position measurement uncertainty

$p_{33}^{1/2} = p_{44}^{1/2} = 200$ m/sec

Model Uncertainty: $\text{Diag} [q_{11}, q_{22}, q_{33}, q_{44}]$

$q_{11}^{1/2} = q_{22}^{1/2} = 0$

$q_{33}^{1/2} = q_{44}^{1/2} = 5$ m/sec

Measurement Noise Uncertainty

$\sigma_x = 30$ m

$\sigma_y = 300$ m

Dynamic Model of Target

$$\dot{\mathbf{x}}_m = \begin{bmatrix} 1 & 0 & T & 0 \\ 0 & 1 & 0 & T \\ 0 & 0 & a_m \cos \theta_m & a_m \sin \theta_m \\ 0 & 0 & -a_m \sin \theta_m & a_m \cos \theta_m \end{bmatrix}$$

$a_m \in \{1\}$

$\theta_m \in \{-30^\circ, 0^\circ, +30^\circ\}$

Temporal Distribution for Births: Poisson with $\lambda_B = 1E-5$

Temporal Distribution for False Alarms: Poisson with $\lambda_{FA} = 4.5E-10$

Probability of Detection: 0.998

Probability of Death: $2E-4$

Probability of No Maneuver: 0.8

Probability of Maneuver: $0.2/n_m$

TABLE 3. HEURISTICS USED

1. Gating: 10
2. Number of global hypotheses retained at each scan: 10
3. Max. number of missed detections permitted for a track: 2
4. Maneuver hypotheses initiated only for well-established tracks for which:

$$p_{33}^{1/2} < 15 \text{ m/sec}$$

$$p_{44}^{1/2} < 15 \text{ m/sec}$$

5. After maneuver is initiated, only transitions permitted are either straight line or same maneuver state.
6. A priori information about target position at birth is ignored, i.e., target position is initialized based on measurement data only.

$$p_{11}^{1/2} = \sigma_x \quad x_1(0) = x_m$$

$$p_{22}^{1/2} = \sigma_y \quad x_2(0) = y_m$$

We have shown in Figs. 8 and 9 the trajectories postulated by the global hypotheses with ranks 2 and 3. It can be seen that the rank 2 global hypothesis is the correct one - it postulates only the correct trajectory. The rank 3 global hypothesis is almost identical to the rank 1 hypothesis. The difference is that the incorrect track is postulated to die at scan 6.




Figure 8. Test Case 1, Rank 2 Global Hypothesis

On examining the remaining global hypotheses which are retained by the algorithm (which are not shown here), we have observed they all postulate the correct

trajectory for the target. Due to the modeling of the target dynamics in discrete time, however, some of them postulate a maneuver initiated at scan 9. Since the position of the target will be influenced by a maneuver only in the next scan, it is only then that the algorithm will reject such incorrect hypotheses. As in the case of global hypotheses with ranks 1 and 3, we have observed that some of the remaining global hypotheses postulate incorrect tracks for short periods of time. Since they are ephemeral, they do not have any adverse effect on the correct target trajectory. This feature of the algorithm, wherein most of the hypotheses that are retained postulate the correct trajectory with some minor differences, illustrates one aspect of the robustness of the suboptimal algorithm.




Figure 9. Test Case 1, Rank 3 Global Hypothesis

Test Case 2

In this test scenario, we simulated two crossing targets along with clutter. The targets cross at the same point in time. At that time one of the targets executes a maneuver too. As in the last scenario, clutter is generated close to targets. The true target trajectories are indicated by continuous lines and the clutter is indicated by []s. This is shown in Fig. 10. The parameters for the simulation and the algorithm are the same as in Test Case 1 with the following addition.

Target 2 - Initial Position
(-9657, -5264)

Speed and Heading
(200 m/s, -45°)

It can be observed that the conditions are particularly severe for the tracking algorithm at the target crossing point where target 1 executes the maneuver. Figure 10 traces the trajectories postulated by the global hypothesis with the highest likelihood (rank 1). Despite the exacting requirements of the scenario, the algorithm identifies the correct hypothesis by the 9th scan. From these two test cases, we see that the algorithm performs very well in spite of the heuristics that have been introduced.




Figure 10. Test Case 2, Rank 1 Global Hypothesis

6. Summary and Conclusions

By formulating the problem of multiobject tracking of maneuvering targets within the framework of hybrid state estimation and using a novel data structure to represent the discrete-state hypotheses, we have constructed the optimal solution. We have provided simulation results which show the feasibility of constructing the optimal algorithm. Since the optimal solution requires exponentially growing storage and computational requirements, we have incorporated suitable hypotheses deletion techniques. Such suboptimal techniques drastically reduce the computational requirements with little loss in accuracy, as has been shown in the simulation results.

The generality of the approach used to construct the algorithm makes it possible to incorporate features other than tracking. For example, the algorithm can be extended to classify targets and identify target measurement characteristics. In fact, the approach used

here defines the framework for implementing algorithms for state estimation in a switching environment.

References

1. Reid, D.B., "An Algorithm for Tracking Multiple Targets," *IEEE Trans. on Automatic Control*, Volume AC-24, Number 6, December 1979, pp. 843-854.
2. Bar-Shalom, Y., "Tracking Methods in a Multitarget Environment," *IEEE Trans. on Automatic Control*, Volume AC-23, August 1978, pp. 618-626.
3. Moose, R.L., "An Adaptive State Estimation Solution to the Maneuvering Target Problem," *IEEE Trans. on Automatic Control*, Volume AC-20, June 1975, pp. 359-362.
4. Willsky, A.S. and H.L. Jones, "A Generalized Likelihood Ratio Approach to the Detection and Estimation of Jumps in Linear Systems," *IEEE Trans. on Automatic Control*, Volume AC-21, February 1976, pp. 108-112.
5. Bar-Shalom, Y. and K. Birmival, "Maneuvering Target Tracking with Variable State Dimension," *ONR Workshop on Command, Control, and Communication*, Silver Springs, Maryland, June 1978.
6. Balakrishnan, A.V., "Development of Minimal-Time Detection Algorithm," *ONR Workshop on Command, Control and Communication*, Monterey, California, August 1982.
7. Kenefic, R.J., "Optimum Tracking of a Maneuvering Target in Clutter," *IEEE Trans. on Automatic Control*, Volume AC-26, Number 3, June 1981, pp. 750-753.
8. Pattipati, K.R. and N.R. Sandell, Jr., "A Unified View of State Estimation in Switching Environments," *American Control Conference*, San Francisco California, pp. 458-465, June 1983.
9. Keverian, K.M. and N.R. Sandell, Jr., "Multiobject Tracking by Adaptive Hypothesis Testing," Report LIDS-R-959, MIT, Cambridge, Massachusetts, December 1979.

**GENERALIZED TRACKER/CLASSIFIER (GTC) - A SYSTEM FOR
TRACKING AND CLASSIFICATION OF MULTIPLE TARGETS BY MULTIPLE SENSORS***

Shozo Mori, Chee-Yee Chong
and Richard P. Wishner

Edison Tse

Advanced Information & Decision Systems
201 San Antonio Cir., Suite 286
Mountain View, CA. 94040

Dept. of Engr.-Econ. Sys.
Stanford University
Stanford, CA. 94305

ABSTRACT

This paper describes a system for tracking and classifying multiple targets using measurements from multiple sensors. This system was implemented based on a general Bayesian theory of multitarget tracking which the authors developed earlier. The related implementation issues in multitarget tracking are also discussed.

1. INTRODUCTION

In [1] and [2] we provided solutions to several problems which had not been solved in the existing multitarget, multisensor tracking literature (surveyed in [3], [4], [5], [6], etc.). The problems were: (a) treatment of non-gaussian target/sensor models, (b) dependence of detection on target states, (c) determination of the likelihood of newly detected targets (how to initiate tracks), and (d) dependence among targets. As indicated in [7], [8], etc., the first problem was less difficult than the others which are, however, obviously important in practice. In the process of solving these problems, we realized that, despite the many algorithms developed in a vast amount of existing literature, there did not exist either a general foundation of the subject or a unified view on it. An attempt was made in [9] and [10] to fit multitarget tracking problems into a subcategory of a special class of dynamical systems.

In [1] and [2], we presented a view which sharply contrasts with this. According to our view, many elements involved in the multitarget tracking problems are in essence random sets as defined in [11], and hence the problems are radically different from the conventional filtering/estimation problems. For example, targets constitute a random set because (in general) the number of targets is unknown, and there is no a priori labelling. The returns from sensors are also random sets because of the random number of the returns and the random ordering. We tentatively call such features of objects random-set nature. Moreover, as will be seen later, many elements (such as tracks and hypotheses) are well understood when they are connected to certain random sets.

The theory on random sets, or stochastic geometry, is mainly concerned with those whose realizations are uncountable sets, e.g., open, closed or convex sets in Euclidean spaces, and is mathematically very sophisticated. Fortunately, in the multitarget tracking problems we are exclusively concerned with random sets whose realizations are finite (or at most countable) sets. Therefore we still can apply standard probabilistic techniques. For example, a finite random set of reals can be probabilistically completely described by specifying the probability on the number N of elements and the joint probability distributions of elements given N 's. However, in order to model the random-set nature, we must require that each of the

*This research was supported by the Defense Advanced Research Projects Agency, Contract MDA-903-81-C-0333.

distributions is interchangeable (permutable). In other words, we can treat a random set as a random element in a disjoint union of direct-product spaces with different dimensions. As a matter of fact, this was the basic approach taken in [1] and [2]. For this reason, the basic objects are generally modeled as a pair consisting of the number-of-elements and a vector with the corresponding dimension.

The main purpose of this paper is to describe a general multitarget tracking system (called Generalized Tracker/Classifier -- GTC) which we implemented based upon our general theory developed in [1] and [2] and to compare this system with other existing algorithms. In the following three sections, we will briefly overview our theoretical results. (One should refer to [1] or [2] for a more rigorous treatment.) Then we discuss the implementation issues and compare the GTC implementation with other algorithms.

2. GENERAL MODEL

A target is a generic name for any object to be tracked and/or classified, and a sensor is a generic name for a device which generates an undetermined number of measurements. In our general model, targets are modeled as one entity, i.e., a pair of the two random elements: a nonnegative random integer N_t representing a constant but unknown number of targets in the "whole world" and a continuous-time Markov process $X(t)$. The Markov process $X(t)$ is on different spaces depending on the number of targets N_t . Therefore, the targets are modeled as an abstract stochastic process $(X(t), N_t)$, called target system state on the countable disjoint union of spaces, which we call target system state space. Each component of the target system state space corresponds to a given number of targets and consists of (is the direct product of) two spaces: the common component space (to which a common state to all the targets in a group may be assigned) and the individual component space. To account for the random-set nature of the targets, the latter space must be an N_t -direct-product of an identical space and the Markov process $X(t)$ must be interchangeable with respect to the ordering (labelling) of targets. Any component of the target system state space is in general a hybrid set, i.e., the direct product of a subset of a Euclidean space (used to represent "continuous" information such as position and velocity) and a finite set (used to represent "discrete" information such as type, maneuvering/operating mode, etc.).

One should not confuse the constant (but unknown) number of targets with that of detected targets, and should note that the constancy assumption is not restrictive but correctly reflects the reality at least conceptually. For example, consider targets whose states reflect their birth-death processes (as modeled in [8]). In reality, the interval between the time when the first target is born and the time when the last target dies is finite. The total number of targets which do exist in subintervals is hence constant

and finite. In a sense, the concept of existence is independent of the concept of time. In practice, however, we may have some difficulties in reflecting a "constant" flow of targets into a surveillance region. In almost all cases, we can resolve this difficulty by choosing appropriate a priori initial distributions of target states.

In general, we assume a multiplicity of sensors, each of which generates sets of measurements (returns) at a finite rate. Each sensor output is again a pair (y, N_M) of two random elements, the number N_M of measurements and a vector y of measurements, and is a random element in a disjoint union of direct-products of identical measurement spaces. Each sensor may however have a different measurement space which is in general, a hybrid set to account for feature-type discrete information as well as continuous-value observations. When a sensor s generates an output (y, N_M) at time t , we call (y, N_M, t, s) a data set and (t, s) a data set index.

Without loss of generality, we can assume that any sensor generates at most one output at a time. One of the key steps in our development is to model a mechanism which relates the target system state (state-of-world) to the sensor outputs (observables). The first assumption is that the sensor source s and the timing t of its output are exactly known and are independent of the targets. Therefore, even when data set indices (t, s) 's are random, we can treat them as constant and known quantities, i.e., all the sensor schedules are predetermined. Let \mathcal{K} be the set of all the data set indices (t, s) 's.

With the standard assumption of no split measurements and no merged measurements, for each $k = (t, s)$ in \mathcal{K} , sensor output (y, N_M) is modeled as follows: We assume a random subset I_D of the target index set $I_T = \{1, \dots, N_T\}$ such that $\#(I_D) \leq N_M$. (In this paper, $\#(E)$ is the number of elements in a set E .) and a random one-to-one function A defined on I_D taking values in the set $J = \{1, \dots, N_M\}$. The random subset I_D is called the (index) set of detected targets and the random function A is called assignment function at k . $i \in I_D$ and $j = A(i)$ means that the i -th target is detected at k and the j -th measurement originates from it. When j in J is not in the image of A , we call the j -th measurement false alarm.

Finally, we assume that every data set is conditionally independent given the target state. With this assumption, we can completely describe the sensor model by, for each (t, s) in \mathcal{K} , specifying (i) the probability of the random subset I_D of I_T given N_T , (ii) the probability distribution of the number N_M of measurements (or equivalently the number N_{FA} of false alarms) and (iii) the probability density function of the measurement vector given I_D , N_M and the random assignment A . All of the above probabilities, distributions and densities are, in general, conditioned by the target system state. In particular, detection of a target may depend on its state. To reflect the random-set nature, each of the conditional probabilities, etc., in the above (i) - (iii) must be interchangeable with respect to the ordering of the targets.

3. GENERAL BAYESIAN FORMULA

Let \mathcal{Z} be the set of all the data sets and \mathcal{K} be the data set index set. Since every sensor generates outputs at a finite rate and it sends at most one data set at a time, \mathcal{K} is at most countable and is one-to-one to \mathcal{Z} . \mathcal{K} can be enumerated in accordance with time in such a way that \mathcal{K} and \mathcal{Z} are totally ordered so that the time components are in a natural order. Thus, for each k in \mathcal{K} , the cumulative data set Z up to k is the collection of all the data sets up to k and

is the total amount of information accumulated at k . Our problem is to calculate the probability distribution of the target system state $(x(t), N_T)$ given the cumulative data set Z up to $k = (t, s)$ in \mathcal{K} , in terms of (i) $P(x(t) | N_T, Z^{(k)})$ and (ii) $P(N_T | Z^{(k)})$. Since the origin of each measurement in each data set is random and there is no a priori labelling, we cannot directly calculate such conditional probability (i) described above. Therefore we must "hypothesize" the origin of each measurement, evaluate each "hypothesis" and calculate the target system state conditional distribution given each hypothesis, which is in fact a common procedure in multitarget tracking. Hypotheses as well as tracks are the most frequently used terminologies but often only vaguely defined. Our definition of hypotheses and tracks is closely related to Morefield's notion in [12] but differs from it in that ours are defined exclusively on the number-of-measurements information.

Let k be a data set index in \mathcal{K} and $J^{(k)}$ be the union of all the sets $\{1, \dots, N_M(k')\} \times \{k'\}$ for all k' up to k where $N_M(k')$ is the number of measurements in the data set indexed by k' . Each element (j, t, s) in $J^{(k)}$ represents the j -th measurement in the data set, generated by sensor s at time t . The random set $J^{(k)}$ is called the cumulative measurement index set up to k . A track at k is a subset of the cumulative measurement index set up to k and a hypothesis at k is a collection of nonempty tracks at k . A track is said to be possible if it contains at most one measurement index in one data set. A hypothesis is said to be possible if it contains only possible tracks and no two tracks in it intersect. Define a random set $\Lambda(k)$ by

$$\Lambda(k) = \{ \{ (A(k')(i), k') | k' \in K(k), i \in I_D(k') \} | i \in I_D^{(k)} \} \quad (1)$$

where $K(k)$ is the set of data set indices prior to and including k , $I_D^{(k)}$ is the set of detected targets in the data set indexed by k , $A(k')$ is the random assignment in the same data set, and $I_D^{(k)}$ is the cumulative set of detected targets up to k , i.e., the union of all the $I_D(k')$ for $k' \in K(k)$. Due to our no-split/no-merge assumption, it is clear that the summation of $\text{Prob.} \{ \Lambda(k) = \lambda | J^{(k)} \}$ over all the possible hypotheses λ at k must be one. When $\lambda = \{ \tau_1, \dots, \tau_n \}$ is a possible hypothesis, in the event $\{ \Lambda(k) = \lambda \}$, there are n targets which are detected at least once up to k and all the measurement indices in each τ_i originate from the i -th detected target.

Since the right-hand side of (1) is a set operation and the random set $I_D^{(k)}$ is not accessible, we must "hypothesize" the origin (in I_T) of each element in $\Lambda(k)$, i.e., a one-to-one random function $\Omega(k)$ from $\Lambda(k)$ into I_T . Therefore, we must calculate (i) $P(N_T | \Lambda(k), Z^{(k)})$, (ii) $P(\Omega(k) | N_T, \Lambda(k), Z^{(k)})$ and (iii) $P(x(t) | \Omega(k), N_T, \Lambda(k), Z^{(k)})$. Due to our interchangeability condition, however, we can show that the probability (ii) is the same for all the possible realization of $\Omega(k)$ and that the state distribution (iii) is invariant under any realization of $\Omega(k)$. Thus, we need only calculate (i) and (iii), which can be done recursively as shown below: Let $k = (t, s)$ be any data set index in \mathcal{K} , (y, N_M, k) the last (current) data set in the cumulative data set $Z = Z^{(k)}$ and $Z = Z^{(k)} = Z \setminus \{(y, N_M, k)\}$, where $\bar{k} = (\bar{t}, \bar{s})$ in \mathcal{K} is the immediate predecessor of k and \setminus is the set subtraction operation. Then, with $\Lambda = \Lambda(k)$ and $\bar{\Lambda} = \{ \{ \tau_i^{(k)} | \tau_i \in \Lambda \setminus \{ \phi \} \} \}$, we have

$$P(\Lambda|Z) = \frac{P(\bar{\Lambda}|\bar{Z}) (N_M - n_D(\Lambda|k))!}{P(Z|\bar{Z}) N_M!} \sum_{N_T=\#(\Lambda)}^{\infty} \frac{(N_T - \#(\bar{\Lambda}))!}{(N_T - \#(\Lambda))!} P(N_T|\Lambda, Z) \mathcal{L}(y, N_M|N_T, \Lambda) \quad (2)$$

where $n_D(\lambda|k)$ is the number of detected targets at k when hypothesized by λ and $\mathcal{L}(y, N_M|N_T, \lambda)$ is the likelihood of the current sensor output (y, N_M) under assumption (N_T, λ) . The actual form of \mathcal{L} is rather complicated and is omitted in this paper (Refer to [1] or [2]). Then, to complete the recursion, the target system state distribution conditioned on $(Z(k), \Lambda(k), \Omega(k))$ must be calculated as well as the conditional probability on N_M according to the updating formulae shown in [1] and [2].

4. INDEPENDENT, IDENTICALLY DISTRIBUTED CASES

Although the general formulation described in the previous section gives us a sound foundation of our theory of multitarget tracking, its implementation may pose a serious problem because evaluation of hypotheses involves all the possible pairs (λ, N_M) . However, with the additional independence assumptions mentioned below, we can reduce the general formula to a form which is more implementationally feasible. We call models which satisfy these assumptions independent and identically distributed (i.i.d.) models. With these assumptions, applicable models are more limited; for example, we cannot allow dependence among targets. However, this class of models is still general enough to include any model in the existing multitarget literature.

First we assume that, given the total number N_T of targets, the targets are represented by a system of i.i.d. Markov processes where the target system state space consists only of the individual target system space. The *a priori* distribution of the number N_M of targets is Poisson with mean ν_0 . The event $\{i \in I_D(k)\}$ in which the i -th target is detected in the data set indexed by $k = (t, s)$ is target-wise conditionally independent and depends only on its own state $X_i(t)$. Thus, we have

$$\text{Prob.}\{i \in I_D(k)|X(t), N_T\} = p_D(X_i(t)|k)$$

with a given detection probability function $p_D(\cdot|k)$. We should note again that $p_D(\cdot|k)$ is a conditional probability so that the probability of the i -th target being detected is described as

$$\text{Prob.}\{i \in I_D(k)|N_T\} = \int p_D(X_i(t)|k) p(dX_i(t)|N_T)$$

which one should not confuse with $p_D(\cdot|k)$ itself.

The number of false alarms in the data set indexed by any k is independent of the target states and of other data sets, and is determined by a given distribution $p_{N_{FA}}(\cdot|k)$ for each k . Given the set $I_D(k)$ of detected targets, the number of false alarms (hence, the number $N_{FA}(k)$ of measurements) and the assignment function $A(k)$, the system of measurements values $y = (y_j)_{j \in \{1, \dots, N_M\}}$ is target-wise conditionally independent and, in particular, the system of false alarm measurements is i.i.d. with a common density $p_{FA}(\cdot|k)$. When $j = A(k)(i)$, the probability density of y_i is given by a usual state-measurement transition probability density function $p_M(y_j|X_i, k)$, i.e.,

$$\text{Prob.}\{y_j \in dy | j = A(k)(i), X_i(t)\} = p_M(y_j|X_i(t), k) \mu_s(dy)$$

where μ_s is an appropriate measure on the measurement value space of sensor s .

With the above additional assumptions, the general formula (2) in the previous section can be reduced to:

$$P(\Lambda|Z) = C^{-1} P(\bar{\Lambda}|\bar{Z}) L_{FA}(y, N_M|\Lambda, k) \prod_{T \in \Lambda} L(Y(T, k)|T, k) \quad (3)$$

where

$$C = P(Z|\bar{Z}) (N_M!) \exp(\nu(\bar{k})(1 - L(\emptyset|\emptyset, k))) \quad (4)$$

is the normalizing constant,

$$L_{FA}(y, N_M|\Lambda, k) = (N_{FA}!) p_{N_{FA}}(N_{FA}|k) \prod_{j \in J_{FA}} p_{FA}(y_j|k) \quad (5)$$

is the false alarm likelihood with N_{FA} being the number of false alarms and J_{FA} being the set of false alarms, and $Y(T, k)$ is the measurement assigned by (symbolically $Y(T, k) = \emptyset$ if T is not assigned any measurement at k). Each of the remaining factors $L(y|T, k)$ in (3) is generally called the track-to-measurement likelihood. They are defined as:

(i) the likelihood of measurement y originating from a target detected before:

$$L(y|T, k) = \int p_M(y|x, k) p_D(x|k) p(dx|\bar{T}, \bar{Z}, k) \quad (6)$$

(ii) the likelihood of a target detected before being undetected:

$$L(\emptyset|T, k) = \int (1 - p_D(x|k)) p(dx|\bar{T}, \bar{Z}, k) \quad (7)$$

(iii) the likelihood of measurement y originating from a newly detected target:

$$L(y|T, k) = \nu(\bar{k}) \int p_M(y|x, k) p_D(x|k) p(dx|\emptyset, \bar{Z}, k) \quad (8)$$

(iv) the likelihood of an undetected target remaining undetected:

$$L(\emptyset|\emptyset, k) = \int (1 - p_D(x|k)) p(dx|\emptyset, \bar{Z}, k) \quad (9)$$

In (4) and (8), $\nu(\bar{k}) = E(N_T|\bar{\Lambda}, \bar{Z}) - \#(\bar{\Lambda})$ is the expected number of undetected targets up to \bar{k} . In (6) - (9), $p(\cdot|T, \bar{Z}, k)$ is the target state distribution at t given track \bar{T} and cumulative data set \bar{Z} .

The updating from $\nu(\bar{k})$ to $\nu(k)$ and from $p(\cdot|\bar{T}, \bar{Z}, k)$ to $p(\cdot|T, Z, k)$ can be done by $\nu(k) = L(\emptyset|\emptyset, k)\nu(\bar{k})$ and

$$p(dx|T, Z, k) = \begin{cases} d^{-1} p_M(y|x, k) p_D(x|k) p(dx|\bar{T}, \bar{Z}, k) & \text{if } Y(T, k) = y \neq \emptyset \\ d^{-1} (1 - p_D(x|k)) p(dx|\bar{T}, \bar{Z}, k) & \text{otherwise} \end{cases} \quad (10)$$

with normalizing constant d or d' .

When we impose the additional assumption that the distribution $p_{N_{FA}}(\cdot|k)$ of the number of false alarm is Poisson with mean $\nu_{FA}(k)$ for each k in \mathcal{K} , we have the following simple formula by repeatedly applying the recursive formula (3):

$$P(\Lambda|Z) = C^{-1} \prod_{T \in \Lambda} L(T, Z) \quad (11)$$

where $L(\tau|Z)$ is the likelihood of track τ , which can be recursively calculated by

$$L(\tau|Z) = \begin{cases} \nu_0 & \text{if } Z = \emptyset \\ L(\tau|\bar{Z}') L(y|\tau, k') / \beta_{FA}(y|k') & \text{if } Y(\tau, k') \neq y \neq 0 \\ L(\tau|\bar{Z}') L(0|\tau, k') & \text{otherwise} \end{cases} \quad (12)$$

and

$$\beta_{FA}(y|k') = \nu_{FA}(y|k') p_{FA}(y) \quad (13)$$

for each k' , $Z' = Z(k')$ and $\bar{Z}' = Z(k')$ with k' being the immediate predecessor of k' .

5. GENERALIZED TRACKER/CLASSIFIER

We have implemented a general algorithm described in the previous two sections in the form of a system called Generalized Tracker/Classifier (GTC). GTC is intended to be used in performance analysis of multi-target tracking systems by Monte Carlo simulation and to implement all the problem-independent parts of the general algorithm. At the current stage, GTC can only handle models for which all the assumptions made in Section 4 are satisfied in addition to those made in Section 3. Equ. (3) was chosen for the GTC's basic equation. The functions in GTC can be divided into the following three major components:

1. hypothesis generation
2. hypothesis evaluation, and
3. hypothesis management

In hypothesis generation all the tracks and measurements are first cross-referenced and all the likelihood functions appearing Equation (3) are calculated. These functions are tabulated into a table called the track-measurement cross-reference table. Then all the old hypotheses and tracks are expanded to include the new measurement assignments by a measurement-oriented tree expansion technique described in [14]. In hypothesis evaluation, all the new hypotheses are evaluated using Equation (3) and all the new tracks are assigned with the updated target state distributions using Equation (10).

As is well known the number of possible hypotheses and that of possible tracks grow very rapidly as the data sets are accumulated. In fact we found that this growth is worse than exponential in most cases. Therefore, as in any implementation of any multi-hypothesis system, we must have reasonably powerful hypothesis management techniques to control the number of hypotheses and tracks. The hypothesis management techniques incorporated into GTC can be categorized as follows:

- (a) hypothesis pruning
- (b) hypothesis combining
- (c) windowing and
- (d) clustering

Each of the above categories will be discussed in the next section.

Since the data used in GTC are structured and the generally large memory space is required, the choice of programming language and the dynamical memory allocation are crucial issues. The first version of GTC was implemented in SAIL. This version was tested in several different target/sensor models, from very simple ones to complicated models with state-dependent detection probabilities and target classification. Currently, implementations in two other versions in C and LISP are under way.

6. HYPOTHESIS MANAGEMENT TECHNIQUES

Unfortunately there is no hypothesis management theory applicable to general multitarget tracking problems. To date however many hypothesis management techniques have been proposed and tested. In this section we will discuss these techniques and describe what techniques we have chosen for the GTC and why. We will describe these techniques according to the four categories of hypothesis management described in the previous section.

(a) Hypothesis Pruning: The existing pruning techniques are further categorized as (a1) fixed-threshold pruning in which insignificant hypotheses are cut and (a2) fixed-breadth pruning in which the number (breadth) of hypotheses is limited. In [13] fixed-threshold pruning (a1) was proposed in which any hypothesis whose probability falls below a predetermined level is pruned. In this method however no consideration of the computing resource or of the external conditions is made. Typically, the uncertainty of the origin of each measurement is not resolved in the first few data sets and is generally resolved as more data sets are accumulated. In such a case, we may want to keep more hypotheses at earlier stages than later stages. On the other hand, how many hypotheses can be stored is actually determined by the available computer resource. From this point of view fixed-breadth pruning (a2) proposed in [14] makes more sense. In this method the hypotheses are ordered by their probabilities and, for a given fixed breadth N , at most N best hypotheses are kept. In an extreme implementation of this technique N is one, i.e., only the best hypothesis is kept. Such a method is called zero-scan algorithm in [13]. Unfortunately this method loses its rationale when we use any form of clustering described later in this section, since some clusters may need more breadth than others under certain conditions and there is no intelligent way of allocating the resource among clusters.

In order to compensate for these shortcomings, we have introduced a new pruning technique called adaptive-threshold/adaptive-breadth pruning. In this method, all the hypotheses are ordered according to their probabilities and the cumulative probabilities from the best hypothesis are calculated; when the cumulative probability exceeds a predetermined threshold, the remaining low-probability hypotheses are discarded. One can view the fixed-threshold pruning as an adaptive-breadth-pruning and vice versa. This new pruning method, however, possesses overall adaptivity in which the breadth in each cluster is adjusted according to the complexity of the measurement data in the cluster and the computer resource is adequately allocated among clusters.

(b) Hypothesis Combining: When the measurement data are confusing, we may have a large number of similar hypotheses with small probabilities. In such a case, unless we combine similar hypotheses, either of the pruning methods described above could fail resulting in the loss of "important" hypotheses. Since the set of all the possible hypotheses is a partition of a probability space with respect to the uncertainty of the origin of each measurement in the past, to combine hypotheses is actually to coarsen this partition. If the partition becomes too coarse, the performance of the tracker may be greatly degraded. Therefore a common strategy in hypothesis combining is to combine only "similar" hypotheses. There are two well documented hypothesis combining methods: (b1) target-state-oriented combining described in [13] and (b2) measurement-index-oriented combining originally proposed in [16]. In the first method, two hypotheses are similar and hence to be combined if they have an identical number of tracks and each track in a hypothesis

has a unique similar track in the other hypothesis. Two tracks are similar if their state distributions are similar. In [17] a criterion for testing the similarity of two gaussian distributions was proposed. In general cases however the distribution similarity should be carefully considered based on the particular target/sensor models. When two hypotheses are combined, each pair of similar tracks is combined and the state distributions are combined with the hypothesis probabilities as weights.

Measurement-index-oriented combining (b2) is actually a classic method proposed almost a decade ago. In this method, for a given M, two tracks with the same measurement assignment in the most recent M data sets are identified. When tracks are identified in this way, we may have sets of identical hypotheses which can be combined naturally. This method is sometimes called the M-scan or depth-M algorithm. This scheme may not be adequate when the arrival of data sets is irregular or sensor characteristics vary widely from sensor to sensor. Moreover, when two tracks are combined, there is no adequate weighting. For this reason we have chosen the target-state-oriented combining in which the determination of two similar tracks and the actual track combining are performed by a user-provided external routine outside the GTC.

As an extreme example of combining, in the JPDA algorithm described in [15] all the hypotheses are combined at each stage so that only one hypothesis is kept for the next stage. It is clear, however, that this is possible only when we have a priori probability-one hypothesis with a priori tracks and there is no newly detected target.

(c) Windowing: When a track state distribution has a reasonable variance and measurement errors are not exceptionally large with respect to the field-of-view of a sensor, one can expect the track-measurement likelihood defined by (6) to be very small except for a limited region. Windowing techniques are generally designed to set an appropriate threshold on the likelihood function so that its value outside this region is considered zero rather than a very small but still positive number, thereby eliminating unnecessary hypothesis expansion. In other words, windowing is screening of data to determine which ones we should try to associate to a given track. For this reason, windowing is often called data validation and, for each track, the set of measurements in which the likelihood function is not zero is called the validation region. When the track state distribution and the measurement error are both gaussian, threshold windowing by the gaussian likelihood functions, corresponds exactly to classical χ^2 - or extended χ^2 -test.

One also may view windowing as immediate pruning in which a branch is cut solely based on one likelihood function. However, since hypothesis evaluation cannot be completed until all the hypotheses are expanded and their probabilities are obtained by normalization, any intermediate pruning may degrade the tracking performance. Therefore, one should choose adequate windowing procedure rather by carefully examining the physical nature of the particular target/sensor models used in the system. For this reason, GTC relies on a user-provided external routine for performing adequate windowing and simply receives the zero likelihood when a measurement data is not validated.

(c) Clustering: Because of our i.i.d. assumptions made in Section 4, when the validation region of each track is not exceptionally large, we can further decompose the basic equation (3) or (11). In such a case, we can group tracks and measurements so that local hypotheses on such a group, called a cluster, may be generated and evaluated locally. For this

decomposition to be valid, any two tracks across two different clusters should not intersect. If this condition is satisfied, the set of global hypotheses can be reconstructed by forming all the possible unions of local hypotheses. Thus clustering techniques are in general methods for maintaining this non-intersection condition. An algorithm for performing such a task was described in [13] and the current GTC adopts this technique.

7. COMPARISON WITH OTHER ALGORITHMS

The most significant difference of our basic formulation from others is that, for each data set, we do not assume the probability of detection but rather specify a detection function which is actually the conditional probability given a target state, i.e., state-dependent detection probability. In many cases, we cannot correctly model sensors without state-dependent detection probability. For example, a target may be out of a sensor field-of-view or hidden otherwise in an unobservable region. If we cannot have a common probability of detection for all the tracks in a data set, the most commonly used binomial distribution on the number of detected targets (originally used in [13] and subsequently used in many others such as [6], [10], etc.) is no longer adequate. Instead the Bayesian expansion via the random set L_D of detected targets and the random assignment function A at each data set k must be employed.

In many applications, however, the detection probability function is determined only by a sensor field-of-view as $p_D(x) = p_{Dmax} \pi_D(x)$, e.g., where

$$\pi_D(x) = \int_{Q_j} g(y - Hx; R) dy.$$

$g(\cdot; R)$ is the zero-mean multi-dimensional gaussian density with variance R, Q_j is the field-of-view and $y = Hx + \text{noise}$ is the measurement equation with error variance R. In such a case, the target-state-measurement transition probability density becomes

$$p_M(y|x) = g(y - Hx; R) / \pi_D(x).$$

on Q_j . Therefore, in such special cases, if a track has a priori state distribution which is gaussian with mean \hat{x} and variance P, the track-measurement likelihood defined by (6) becomes $p_{Dmax} g(y - H\hat{x}; HPH^T + R)$. Moreover the track-no-measurement likelihood,

$$\int (1 - p_D(x)) g(x - \hat{x}; P) dx$$

can be approximated by $(1 - p_{Dmax})$ if Hx is well inside Q_j and $HPH^T + R$ is not exceptionally large. Exactly these two terms appeared in Reid's algorithm in [13].

The crucial difference of our algorithm in such gaussian cases is the treatment of likelihood $L(y, \tau)$ where τ is formed for the first time. In Reid's algorithm, such an $L(y, \tau)$ is constant and is called new target density p_{NT} . For example, when we observe targets moving at almost constant velocities in a one-dimensional space with a relatively high detection probability p_{Dmax} , as data sets are accumulated we expect newly detected targets to appear only on the edges of the field-of-view. Only by calculating this likelihood as a function of measurement value, can we adequately incorporate such an effect into track initiation. In many applications, the processing of several measurements right after a track is formed for the first time

is very important since it serves as a multi-scan target detector. In such cases, an exact calculation of this likelihood may be very important.

As shown in [13], in a typical situation, we must form many hypotheses in several data sets after a track is formed for the first time and these hypotheses eventually become a single probability-one hypothesis by means of pruning and combining. After a track is "confirmed" or "initiated" in such a way, the GTC algorithm behaves in a way very similar to that of JPDA algorithm described in [15] if the likelihood of receiving signals from newly detected targets is very low in the validation regions of the track. As a matter of fact, if we set the newly detected target likelihood functions to be all zero and start GTC algorithm with a single a priori probability-one hypothesis, we can reduce our algorithm to JPDA with gaussian assumptions and appropriate approximations.

Equ. (11) in Section 5 gives a general formula both for the track-likelihood approach and for the batch-processing approach. An algorithm using the former approach is described in [14] while one using the latter is described in [12], both with gaussian models.

8. CONCLUSIONS

A system for the tracking and classification of multiple tracks by multiple sensors, called GTC, has been described. This system was built as the first implementation of the general Bayesian multitarget tracking formulation which the authors had developed earlier. By this implementation, we can reasonably handle the problems of state-dependent detection probability and target initiation processes. Also, within the i.i.d. assumptions, fairly general target/sensor models can be handled. As seen in Section 4, however, nonlinear filtering problem is always a major subproblem in any multitarget tracking problem. Beyond Kalman or extended Kalman filters or totally discrete estimation problems, nonlinear filtering problems are very difficult to handle. Thus good solutions to such filtering problems are still essential to a successful implementation when non-gaussian models are required. Although the problem of dependence among targets has been theoretically solved, further studies are necessary before a general algorithm capable of handling such complicated situations can be implemented.

REFERENCES

- [1] S. Mori, C.Y. Chong, E. Tse and R.P. Wishner, "Multitarget Multisensor Tracking Problems - Part I: A general Solution and a Unified View on Bayesian Approaches," submitted to IEEE Trans. Automat. Contr.
- [2] S. Mori, C.Y. Chong, E. Tse and R.P. Wishner, "Multi-target Multi-sensor Tracking Problem: A General Bayesian Approach," Proc. of 1983 American Control Conference, pp. 452-457, San Francisco, June, 1983.
- [3] Y. Bar-Shalom, "Tracking Methods in a Multitarget Environment," IEEE Trans. Automat. Contr., vol. AC-23, pp. 618-626, Aug. 1978.
- [4] H.L. Wiener, W.W. Willman, I.R. Goodman and J.H. Kullback, "Naval Ocean-Surveillance Correlation Handbook, 1978," NRL Report 8340, Naval Research Lab., Washington, D.C., 1979.
- [5] I.R. Goodman, H.L. Wiener, and W.W. Willman, "Naval Ocean Surveillance Correlation Handbook, 1979," NRL Report 8402, Naval Research Lab., Washington, D.C., 1980.
- [6] I.R. Goodman, "A General Model for Contact Correlation Problem," NRL Report 8417, Naval Research Lab., Washington, D.C., 1983.
- [7] D.B. Reid, "A Non-Gaussian Filter for Tracking Targets Moving over Terrain," Conf. Record of 12th Asilomar Conf. on Circuit, Systems and Computers, Nov., 1978.
- [8] I.R. Goodman, "A General Model for the Multiple Target Correlation and Tracking Problem," Proc. of 18th IEEE Conf. on Decision and Control, Fort Lauderdale, FL., Dec. 1979.
- [9] K.R. Pattipati, N.R. Sandell, Jr. and L.C. Kramer, "A Unified View of Multi-Object Tracking," Proc. 4th MIT/ONR Workshop on Distributed Information and Decision Systems Motivated by Command-Control-Communication Problems, vol.1, pp. 115-135, San Diego, Ca., June 1981.
- [10] K.R. Pattipati and N.R. Sandell, Jr., "A Unified View of State Estimation in Switching Environments," Proc. of 1983 American Control Conference, pp. 458 - 465, San Francisco, June, 1983.
- [11] G. Matheron, Random Sets and Integral Geometry, John Wiley & Sons, Inc., 1975.
- [12] C.L. Morefield, "Application of 0-1 Integer Programming to Multi-Target Tracking Problems," IEEE Trans. Automat. Contr., vol. AC-22, no. 3, pp. 302-312, June 1977.
- [13] D.B. Reid, "An Algorithm for Tracking Multiple Targets," IEEE Trans. Automat. Contr., vol. AC-24, No. 6, pp. 843-854, Dec. 1979.
- [14] K.M. Keverian and N.R. Sandell, Jr., "Multi-Object Tracking by Adaptive Hypothesis Testing," Report LIDS-R-959, MIT, Dec. 1979.
- [15] T. Fortmann, Y. Bar-Shalom and M. Scheffe, "Joint Probabilistic Data Association for Multiple Targets in Clutter," Proc. 19th IEEE Conf. Decision and Contr., Albuquerque, NM., Dec. 1980.
- [16] R.A. Singer, R.G. Sea and K. Housewright, "Derivation and Evaluation of Improved Tracking Filters for Use in Dense Multitarget Environments," IEEE Trans. Inform. Theory, vol. IT-20, pp. 423 - 432, July, 1974.
- [17] R.A. Reid, "A Multiple Hypothesis Filter for Tracking Multiple Targets in a Cluttered Environment," Report LMSC-D560254, Lockheed Missiles & Space Co., Inc., Sunnyvale, Ca., Sept. 1977.

A FRAMEWORK FOR EVALUATION OF SURVEILLANCE ALGORITHMS*

✓ Mr. Jerold L. Weiss
ALPHATECH, Inc.
111 Middlesex Turnpike
Burlington, Massachusetts 01803

R.
Professor Robert Tenney
Massachusetts Institute of Technology

I. Introduction

The special structure of surveillance systems has been recognized and exploited in the development of relevant mathematical models for some time. (e.g., see [1], [2]). These models provide the necessary relationships for applying probabilistic techniques to surveillance problems and therefore determine the complexity of the resulting solutions. Typically, such solutions are very complex and require unreasonably large amounts of memory. For this reason, various algorithms have been developed which restrict the amount of complexity at the possible expense of system performance. The purpose of this paper is to formulate a mathematical framework in which analytical comparison of these algorithms may be performed.

The special structure of surveillance systems may be seen by examining two important and related surveillance problems; target tracking and identification. In the most general setting, identification may include tasks such as data association, data type identification, target type identification and maneuver detection. These tasks are "discrete" types of tasks in that a decision amongst a finite number of hypotheses must be made. Tracking problems on the other hand, are "continuous" in nature since the position/velocity of each target may take on any of a continuum of values. The result of these observations leads to the fact that the mathematical models which characterize the unknowns in a surveillance system contain both discrete and continuous states. Note that the measurements which provide information for performing tracking and identification tasks may be discrete or continuous.

Systems which require such a "hybrid state" formulation are not necessarily unreasonably complex. However, in addition to this hybrid structure, we must also consider that these problems are dynamic. In particular, the discrete states which characterize uncertainty in identification problems, may change with time. This behavior reflects, for example the arbitrary ordering of radar data, the scintillation of clutter returns, and the execution of maneuvers by the target. As we will see, the complexity of certain solutions to surveillance problems is directly related to the dynamic behavior of the discrete states.

In contrast to the resultant complexity associated with dynamic hybrid state space characterizations, many problems in surveillance can be simplified by recognizing and modelling their hierarchical structure. That is, it is frequently advisable to consider a subset of hybrid state problems in which the discrete state, or states, evolve independently in time and the continuous states and observations are explicitly dependent on the dynamic realization of the discrete variables.

In the following sections we will review a widely used mathematical model and the methods by which sufficient statistics for estimating the continuous states (tracking) and for making decisions about the discrete states (identification) are obtained. Algorithms which

attempt to reduce the computational complexity of these calculations are then considered in a framework which allows them to be compared. The basis for this framework is the view that each algorithm is in effect, computing approximations to the desired statistics.

II. Mathematical Foundations

A precise description of the dynamic evolution of states in a surveillance system is given by the following mathematical model,

$$x_{k+1} = F(y_k)x_k + w_k \quad (1)$$

$$y_k = H(y_k)x_k + v_k \quad (2)$$

where at time k , the continuous state of the system is x_k , the discrete state is y_k , the observable quantity is y_k , and w_k and v_k are white noises whose covariance matrices, Q_k and R_k , in general may also depend on y_k (i.e., $Q_k = Q(k, y_k)$, $R_k = R(k, y_k)$). Note that the case where F and H are time varying and a control input is present in (1) may be easily included and the results that follow will remain unchanged.

The parameter, y_k , is called the status parameter and takes on one of a finite number of values in a set S_k at each time k . We model the hierarchical structure mentioned previously by assuming that y_k is a Markov chain with known one step transition probabilities. That is,

$$p(y_k = s_i) = \sum_{j=1}^L \Pi_{ij} \cdot p(y_{k-1} = s_j) \quad (3)$$

where $S_k = \{s_1, s_2, \dots, s_L\}$ and $\Pi_{ij} = p(y_k = s_i | y_{k-1} = s_j)$. The problems of identification and tracking in surveillance systems reduces, in this framework, to the estimation of y_k and x_k based on a set of observations, $Y^m = \{y_i : 1 \leq i \leq m\}$. As is the case with all dynamic estimation problems, we can consider smoothing ($k < m$), filtering ($k = m$) and prediction ($k > m$) problems depending on the goals of the specific surveillance system. In each of these cases, estimates may be computed if certain sufficient statistics are known.

For example, if estimates of x_k are of primary importance, then we desire a characterization of the conditional probability density function (pdf), $p(x_k | Y^m)$. Other statistics of importance to identification and tracking include $p(y_k | Y^m)$ and $p(x_k, y_k | Y^m)$ and $p(y_k | y^k)$. For simplicity, in all that follows we will concentrate on characterization of the marginal pdf's for x_k and y_k .

The solution to the problem of characterizing the marginal pdf's is obtained by a direct application of the partitioning theorem [3].

$$p(x_k | y^k) = \sum_{r^k} p(r^k | y^k) p(x_k | y^k, r^k) \quad (4)$$

where $\{r^k\} = \{r_j^k; 1 \leq j \leq k\}$. This result is very useful for two important reasons. First, the conditional pdf's, $p(x_k|Y^k, r^k)$ are Gaussian (under the assumption that x_0 is Gaussian) and are computable from Kalman filter equations matched to the various status sequences [4]. Secondly, the status sequence probabilities can also be computed recursively. In particular, these equations separate into predict and update operations and are given by,

$$\text{Predict: } p(r^{k+1}|Y^k) = p(Y_{k+1}|Y^k)p(r^k|Y^k) \quad (5)$$

$$\text{Update: } p(r^k|Y^k) = \frac{p(Y_k|r^k, Y^{k-1})p(r^k|Y^{k-1})}{p(Y_k|Y^{k-1})} \quad (6)$$

Finally we have,

$$p(Y_k|Y^k) = \sum_{r^{k-1}} p(r^k|Y^k) \quad (7)$$

What the partitioning approach has done is allow us to characterize the desired pdf's by a finite, albeit growing, number of parameters. In effect, it has increased the number of hypotheses which must be considered at time k from the set of values which Y_k may realize, to the set of values which r^k may realize. This growing number of hypotheses with time presents the fundamental limitation to solving surveillance problems via pdf computation and provides the primary motivation for considering suboptimal (i.e., approximate) algorithms.

The most commonly used algorithms in practical surveillance systems can all be viewed as methods of approximating the desired pdf's which are optimally computed from (4) through (7). These approximations involve some form of hypothesis reduction technique in which the number of status sequences, r^k , which must be considered at each time, is kept under some limit. Table 2-1 lists the most common types of hypothesis reduction techniques and is organized to show that these techniques can be classified as either pruning algorithms or merging algorithms.

1. Pruning Algorithms

- Screening (Gating)
- Decision Directed (Hypothesis Testing)
- Maximum Likelihood

2. Merging Algorithms

- Local Moment Invariance
 - fixed hypothesis selection
 - adaptive hypothesis selection
- Best Local Gaussian Representation

Table 2-1. List of Hypothesis Reduction Techniques

III. Comparison of Hypothesis Reduction Techniques

Comparing the various hypothesis techniques in a particular surveillance system usually involves extensive Monte Carlo simulation and can be quite costly. In addition, such simulation methods provide little insight into the importance of various parameters of the algorithms. In this section, a framework and some sample results are developed which illustrate the advantage of using analytical methods in the comparison and evaluation of surveillance algorithms.

The hypothesis reduction techniques of Table 2-1 are used in surveillance algorithms in a variety of

ways [5]. In order to compare the effectiveness of each of these techniques we will consider a class of fixed structure algorithms which differ only in the type of hypothesis reduction techniques that are used. The algorithm is executed as follows:

1. Compute the desired pdf. [e.g., $p(x_k|Y^k)$] up to some maximum time, $k=N$, at which point some form of hypothesis reduction technique is necessary to reduce the computational burden.

2. Apply a single hypothesis reduction technique and interpret the remaining hypotheses as an approximate pdf. (e.g., if at time $k=N$, there are 2^N hypotheses corresponding to each r^N sequence when $S = \{0, 1\}$, then the actual pdf $p(x_k|Y^k)$ is a weighted sum of 2^N Gaussian densities; Eq. (4). If we prune half of the hypotheses and renormalize their posterior probabilities, then the pdf approximation is a weighted sum of 2^{N-1} Gaussians).

3. Compute an approximate pdf at time $k=N+1$ by applying the recursive pdf equations to the remaining hypotheses as if they represented the actual pdf.

4. Repeat Step 2 and Step 3.

Figure 1 shows the optimal algorithm and the class approximate algorithms described above in block diagram form. The predict and update equations of the Kalman filters associated with the computation of $p(x_k|r^k, Y^k)$ and the predict and update equations given by Eqs. (5) and (6) are shown explicitly at each time step. The hypothesis reduction technique is contained in the approximation block and is executed at each discrete time k .

As the figure suggests, the optimal estimates (of x_k or Y_k or both) can be obtained only when the actual pdf is available. The class of algorithms described above is capable of providing suboptimal estimates based on approximate pdfs which, apparently become less accurate with time. The comparison of various hypothesis reduction techniques is now accomplished by quantifying the accuracy of the pdf approximations which result from specific techniques and comparing the results.

The class of probability density functions with which we are dealing constitutes an abstract vector space. The predict, update and approximation algorithms can be viewed as operations which map a pdf at time k , $p(k)$, into $p_p(k)$, $p_u(k)$, and $\hat{p}(k)$ respectively. The accuracy with which the pdf $\hat{p}(k)$ approximates $p(k)$ can be measured by considering a vector distance measure, $D(\hat{p}(k), p(k))$ defined on the vector space of pdf's. Many such measures are available and a few are defined below.

The L_2 norm has many useful properties and is defined by;

$$L_2[\hat{p}, p] = \int_X (\hat{p}(x) - p(x))^2 dx \quad (8)$$

where X is the state space for the random variable, x . Note that both marginal pdf's, $p(x_k|Y^k)$ and $p(Y_k|Y^k)$ as well as the joint pdf $p(x_k, Y_k|Y^k)$ can be considered in (8) by interpreting the integral in the Stieltjes sense.

The Bhattacharyya distance is defined as the negative log of the B-coefficient, ρ , which is defined by,

$$\rho[\hat{p}, p] = \int [\hat{p}(x)p(x)]^{1/2} dx \quad (9)$$

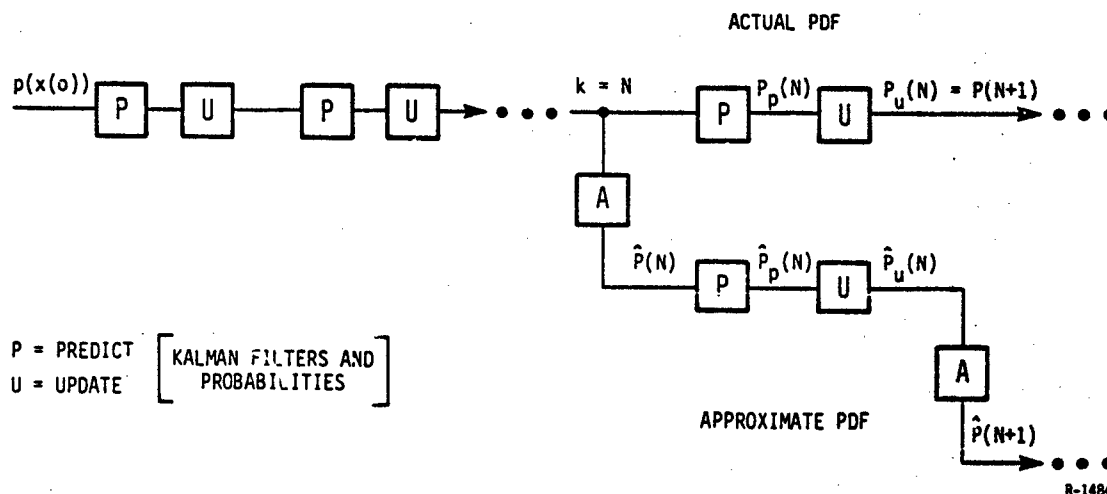


Figure 1. Visualization of PDF Computations

Another useful measure is the Kolmogorov distance or L_1 norm which is defined by,

$$K[p, p] = \frac{1}{2} \int |p(x) - p(x)| dx \quad (10)$$

Ultimately, we would like to be able to choose one of these measures and solve for the distance between $\hat{p}(x)$ and $p(x)$ at each time k . This distance must, clearly, be a function of the mathematical model of Eqs. (1) and (2) and the specific approximation which is used. It is assumed that hypothesis reduction techniques which result in better pdf approximations (smaller distance measures) should result in better surveillance system performance.

As an example of the types of results which are useful, we now consider the K-distance as a measure of performance and state three theorems relating the measure $K[\hat{p}(k), p(k)]$ to $K[\hat{p}(k+1), p(k+1)]$. The proofs of these theorems can be found in reference [6]. (See Figure 1 for definitions).

Theorem 1.

Given $K[p(k), p(k)] = \delta$, then after the prediction operation,

$$K[p_p(k), p_p(k)] < \delta$$

Theorem 2.

Given $K[p_p(k), p_p(k)] = \delta$, then after an update operation,

$$K[p_u(k), p_u(k)] < 2\delta$$

Finally, since K is a valid norm (it obeys the triangle inequality) we can state our last theorem.

Theorem 3.

Let the K-distance across the approximation operation at time k be denoted by δ_k . Then

$$K[p(k+1), p(k+1)] < 2\delta_k + \delta_{k+1}$$

The results above suggest that the distance across the approximation segment are the fundamental determinants of the performance capabilities of a surveillance algorithm which utilizes a specific hypothesis reduction technique. In the next section, two widely used reduction techniques are compared analytically and by simulation.

IV. Comparison of Maximum Likelihood Pruning and Local Moment Invariance Merging Algorithms

As we have stressed before, hypothesis reduction techniques can be viewed as pdf approximation methods. As a simple example of this concept, consider a pdf for some random variable, x , say,

$$F(x) = p_1 f_1(x) + p_2 f_2(x) \quad (11)$$

where $p_1 + p_2 = 1$, $f_1(x)$ are normal densities with means m_1 , and variances σ_1^2 respectively. Thus, x might represent an element of the continuous random vector x_k in (1) and (2), and $F(x)$ may represent an approximation to the posterior pdf which, in general, is a weighted sum of many Gaussians. Two approximations to $F(x)$ are defined below;

$$F_p(x) = f \arg \max_i p_i(x) \quad (12)$$

$$F_m(x) = N_x[m = p_1 m_1 + p_2 m_2; p_1(\sigma_1^2 + m_1^2) + p_2(\sigma_2^2 + m_2^2) - m^2] \quad (13)$$


where the notation $N_x[m; \sigma^2]$ denotes a Gaussian density, in x , with mean m , and variance σ^2 .

The function $F_p(x)$ in (12) is the pdf approximation generated by a Maximum Likelihood Pruning rule for hypothesis reduction. That is, if we are considering two hypotheses, each of which give rise to a Gaussian density for x and their corresponding posterior probabilities p_i , then the pruning rule simply drops the parameters associated with the smallest posterior probability.

On the other hand, the function $F_m(x)$, is the pdf approximation associated with a certain type of hypothesis merging algorithm. Again, if we view p_i as a posterior hypothesis probability and $f_i(x)$ as an hypothesis conditioned pdf of x , then the merging rule

places the two hypotheses with a single one. The single hypothesis conditioned pdf is Gaussian with the first two moments being equal to the first two moments of the weighted sum of two Gaussians defined by (11). Hence the name, local moment invariant merging).

Now, as a result of the reasoning in the last section we can compare these two hypothesis reduction techniques by examining the behavior of the K-distance as a function of the specific densities involved. Fig. 1 shows the numerically computed value of $K[F_p, F]$ and $K[F_m, F]$ when $\sigma_1^2 = \sigma_2^2 = 1$, $p_1 = p_2 = 0.5$, $m_1 = 0$, and m_2 varied. It is quite clear that, for this particular set of pdf's, neither pruning nor merging are superior methods in all circumstances. Pruning seems to be a superior approximation when m_2 is large, while merging is far superior, in terms of K-distance, when m_2 is small. Thus, in practical surveillance systems, it is important to consider the difference between the means of pdf's corresponding to two hypotheses before deciding whether to merge the two hypotheses or prune the most likely one.



$$\begin{aligned} F(x) &= 1/2 [f_1(x) + f_2(x)] \\ f_1(x) &= N_x(0;1) \\ f_2(x) &= N_x(m_2;1) \\ F_p(x) &= f_1(x) \text{ or } f_2(x) \\ F_m(x) &= N_x[m_2/2; 1:m_2^2/4] \end{aligned}$$

Figure 2

K - Distance Between Actual and Approx. PDF's

Similar affects are seen in the performance of these two algorithms in a realistic scenario. Fig. 3 shows the results of a detailed simulation of the two hypothesis reduction techniques (pruning and merging) on a maneuvering target detection/estimation problem. The general algorithm of section III using these two specific reduction methods was applied to the problem of estimating the position of a target whose planar dynamics could be described by one of three possible linear difference equations. Using noisy measurements of position, each estimator was simulated over an initial time interval where the target moves in a straight line with constant velocity followed at time $k=30$, by either a left or a right turn of constant angular velocity. The performance of each algorithm is assessed by examining the RMS position estimation error vs. time. (More details of the algorithms can be found in [6].)

When the actual pdf is computed up to time $k=N=2$, and subsequent approximations keep the number of hypotheses reduced to $3^N=9$, we see that after the onset of a maneuver at $k=30$, the merging technique produces a position estimate whose RMS error is significantly

COMPARISON OF ALGORITHMS N=2





Figure 3. Comparison of Pruning (PRN) and Merging (MRG) Hypothesis Reduction Techniques With Max Number of Hypothesis of 3^N ($N=2$ and $N=3$)

smaller than that obtained by pruning. Furthermore, the pruning algorithm only performs about as well as a Kalman filter (KF) based on the constant velocity straight line motion model with an artificially increased design process noise covariance. However, when $N=3$, so that $3^N=27$ hypotheses are retained, then it is the pruning algorithm whose RMS estimation error is lower following the onset of the maneuver.

Thus, we have seen analytically by comparing K-distances, and experimentally by simulating specific algorithms, that neither maximum likelihood pruning nor local moment invariance merging result in superior surveillance algorithms in every instance.

V. Conclusions

By viewing the hypothesis reduction techniques which appear in many surveillance algorithms as pdf approximation methods we have developed an analytical framework in which these algorithms may be compared. It was concluded that the quality of any surveillance algorithm can be measured by the distance between the pdf approximations before and after the hypothesis reduction technique is implemented. By examining a simple example of two commonly used reduction techniques, we were able to show that neither technique is always superior and could point to some of the parameters which might influence a decision about which technique to employ. In addition, simulation of the two algorithms in a planar maneuvering target tracking problem exhibited the same affects that were pointed out by the simple analytical results.

Besides the results described above, a number of important issues can be addressed within the framework that has been described. First, since the distance measures described in Section III are, typically, difficult to compute, it is advantageous to consider computable bounds on these distances. Some work in this area has been described in [6] and [7] although these bounds can be very loose [6]. Better (tighter) bounds are necessary if accurate comparisons are to be made and/or if these bounds are used in adaptive algorithms. Another issue which requires some attention is the relationship of distance measures to more relevant figures of merit like RMS estimation error. Finally, since we are comparing surveillance algorithms by examining their relationship to the optimal algorithm which would compute the desired pdf's exactly, the issue of how well this optimal algorithm can perform is of fundamental interest. Ultimately, there is always a fundamental limit on the quality of any surveillance system. This limit is a function of the parameters which describe the underlying phenomena (e.g., signal to noise ratios, likelihood of discrete state changes, continuous state observability, etc.) and the performance of the optimal algorithm is one means of assessing the important parameters of these systems.

Acknowledgement

The author would like to thank Dr. Trevini Upadhyay of The Charles Stark Draper Laboratory for his considerable support and insight during the course of the thesis research which led to these results.

References

1. Alspach, D.L., "A Gaussian Sum Approach to the Multi-Target Identification - Tracking Problem", Automatica, Vol. 11, pp. 285-296, 1975.
2. Reid, D.B., "An Algorithm for Tracking Multiple Targets", IEEE Trans. On Automatic Control, Vol. AC-24, No. 6, Dec. 1979.
3. Lainiotis, D.G., "Optimal Adaptive Estimation: Structure and Parameter Adaptation", IEEE Trans. On Automatic Control, V 16, No. 2, April 1971.
4. Ackerson, G.A., and Fu, R.S., "On State Estimation in Switching Environments", IEEE Trans. On Automatic Control, V15, 1970.
5. Kurien, T., Blitz, A., and Washburn, R., "Optimal Maneuver Detection and Estimation in Multiobject Tracking", ALPHATECH Technical Paper TP-182, ALPHATECH, Inc., Burlington, Mass., Sept. 1983.
6. Weiss, J.L., "Comparison of Finite Filtering Methods for Status Directed Processes", S.M. Thesis, M.I.T., Cambridge, Mass., 1983.
7. Lainiotis, D.G., and Park, S.K., "On Joint Detection, Estimation and System Identification: Discrete Data Case", Int. Journal of Control, 1973, Vol. 17, No. 3, p. 609.

INCA: AN ENVIRONMENT FOR EXPLORATORY DEVELOPMENT OF TACTICAL DATA FUSION TECHNIQUES

LEONARD S. GROSS

HAROLD J. PAYNE

VERAC, INCORPORATED, 10975 Torreyana Rd., Suite 300
San Diego, California 92121

ABSTRACT

The Intelligent Correlation Agent (INCA) project, sponsored by DARPA, is directed at the development and demonstration of advanced technologies for supporting data fusion. A broad spectrum of technologies, including Artificial Intelligence, emerging personal workstations, and statistical decision theory, have been examined and used in this project. Central to recent work has been the use of symbolic computers, using object-oriented/messagepassing techniques. The result has been an environment which elevates the level at which developers and analysts can explore fusion techniques. Powerful graphics interaction has been so successful that the current INCA system has been used as a prototype workstation for operational data fusion systems. An incremental development methodology, using feedback from experienced analysts, is being employed in the system development. This paper discusses several of the key aspects of the INCA system.

INTRODUCTION

The Intelligent Correlation Agent (INCA) project, sponsored by DARPA, is directed at the development and demonstration of advanced technologies for supporting tactical data fusion. A broad spectrum of technologies, including Artificial Intelligence, emerging personal workstations, and statistical decision theory, have been examined and used in this project. Central to recent work has been the use of symbolic computers, using object-oriented/message-passing techniques. The result has been an environment which elevates the level at which developers and analysts can explore fusion methods. Powerful graphics interaction has been so successful that the current INCA system has been used as a prototype workstation for operational data fusion systems. An incremental development methodology, using feedback from the experienced analysts, is being employed in the system development. This paper discusses several of the key aspects of the INCA system.

The first section of this paper describes some of the research goals that initially motivated the work. Then a section describing the novel software/hardware aspects of this project is presented. In the third section, we focus on the current machine functionality in both the algorithmic development and workstation prototyping. Next, a discussion of how expert assistance is being integrated into INCA is presented. Finally, future plans for the project are discussed.

INITIAL MOTIVATION FOR THE INCA WORK

Our project focus is the generation of enhanced targeting information by integrating data from multiple sources. In particular, we wish to address the problem of creating ocean surveillance scenes where the data rate per platform is often quite low and the content of the data does not result in unique identification of the platform. As a consequence, there exists consider-

able uncertainty in the association of reports to specific platforms, resulting in ambiguity in the ocean surveillance scene (i.e., the set of platform tracks with uniquely associated reports). Of course, sensor systems are continually supplying new data and this must be incorporated into the evolving scenario.

The basic approach taken to resolve the ambiguity problem is shown in Figure 1. At the top of the tree is the set of initial data from which a set of candidate scenes are generated. Since there is confusion in the data, several scenes may be viable alternatives. The next partition of data (for example, the newest set of data to arrive) is then used to extend some number of scenes from the previous cycle. The obvious combinatorial explosion must somehow be controlled.




Figure 1. INCA: Intelligent Correlation Agent Explosive Growth of Alternative Scenes

An immediate question to address is what type of fusion process is occurring at each node. In general, there are a large number of fusion techniques available. However, it is often necessary to guide any such algorithms to a successful or computationally acceptable solution. At the beginning of the project the knowledge of when to employ, or how to tune, an algorithm was not fully understood. Thus, an initial goal for INCA was the building of a system in which a high degree of user interaction would allow analysts to develop "rules" about when and how to apply different correlation methodologies.

In real-world data processing, human interaction with most processing operations must be limited to produce results in a timely fashion. Hence, the automation of selection and control of fusion node processing was needed. In fact, because of the branching and multiplicity of algorithms to be supported, it was clear that operating the system would require significant expertise. Thus, a secondary thrust of the project was the incorporation of an "expert system" to run or provide assistance to future users.

The project also capitalizing on other artificial intelligence techniques to address the search problems intrinsic in the tree formulation of multiple scene generation.

INCA HARDWARE AND SOFTWARE

Researchers in data fusion are well aware of the difficulty inherent in data fusion processing and the complexity of software development. The use of Fortran code operating on large, time-shared computers is the traditional development tool for testing new techniques and methodologies. The goals for this project lead us to consider new approaches to the correlation software development. Initial work on INCA was undertaken using Interlisp on Dec-10 systems. The power and superior programming environment was recognized immediately as a tool to get leverage over problem complexity. In addition to the obvious advantages of an interactive language with dynamic storage allocation, our motivation for a Lisp-like language was influenced by the need to include limited AI aspects in the project. Additionally, the use of list structures for track and hypothesis representation seemed particularly appealing.

Our initial Interlisp code (1981), which had very weak correlation algorithms, was used in a real-time experiment as an off-line, data quality assurance system.

The success of the initial code reinforced our belief that Lisp-like languages were especially suitable for exploratory research. However, the need for very tight, man-machine coupling was not well supported by the available mainframe graphics. At about the time of the completion of the experiment, a new type of hardware was being brought to realization. Systems based on ARPA/ONR research in artificial intelligence. Generically, these machines are referred to as Lisp machines, since they incorporate special hardware to support the Lisp language. Figure 2 presents some of the key aspects of Lisp machine systems currently available. Our initial hardware, A Symbolics LM-2, was delivered in December of 1981. As significant as the raw power, address space, and superior man-machine interface of these machines, however, is the powerful data abstraction constructs of the language which constitute the Flavor system that is provided as an extension of the Lisp language. The Flavor system allows programmers to develop structures that are closely related to the physical or theoretical objects that the researcher is trying to model. Thus, objects such as "reports" or "sensors" are manipulatable entities in the system. For example, a track "object" may store its reports as a list of the form:

```
(REPORT-OBJECT-1 REPORT-OBJECT-2 ...);
```

adding a report to a track is simply accomplished by the statement

```
(SEND TRACK ':ADD-REPORT REPORT).
```

The Flavor system is a type of object-oriented programming approach, distinguished by the use of "message passing" facilities. Message passing has the very powerful feature that daemons and triggers can be attached to the abstract data types. For example, a daemon can be added to the :ADD-REPORT message so that the track state is automatically recalculated. These capabilities collectively support very flexible implementation schemes.

THE CURRENT INCA SYSTEM

This section will present the tools that are currently available to the analyst for developing an ocean surveillance scene and for exploring new correlation methodologies. This information will be presented by showing several screen images from the INCA system.

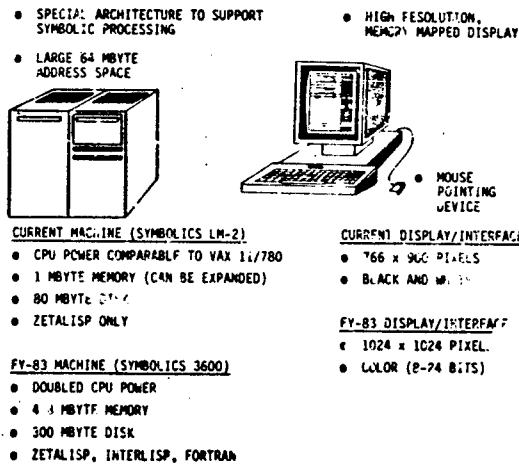


Figure 2. Advanced Personal Computer Hardware

Figure 3 shows the main window of the INCA system. Virtually all user interaction is via a mouse and mousesensitive icons and menus on the screen. The large gridded window at the left is for geographic displays; it can be divided into as many as four separate geographic display panes. The rectangular boxed structure below the INCA logo is called a blackboard and is central to the user interface. The blackboard is divided into levels (rcws), each of which represents a step in the correlation process. Each level of the blackboard is divided horizontally into a set of bins; the logical definition of the bin levels is shown in Table 1. Each of the bins in a level contains a definition of how its data were generated and also contains the result of executing that definition. We have intentionally separated logically distinct aspects of the processing in order to give the operator/analyst the maximum visibility into and control over the processing. Immediately under the blackboard is a set of menus which present many of the top-level functions to the user. These include the functions to control the display, to define and execute a bin, and a set of manual track and report editing commands.

Table 1. Steps in Recursive Scene Generation

ACTIVITY	LEVEL
SELECT PARTITION OF NEW DATA	R
GENERATE FEASIBLE EXTENSIONS TO THE TRACKS IN THE OLD SCENE(S) USING NEW REPORTS	FT
GENERATE A SET OF ALTERNATE WORLD VIEWS, FEASIBLE SCENES	FS
SELECT A SET OF THE FEASIBLE SCENES TO BE EXTENDED ON NEXT ITERATION	S

Figure 4 shows how the user interacts with INCA in order to generate a partition of data on which he wishes to do some processing. The user had moused the DEFINE AND EXECUTE menu item followed by a selection of the first bin at the R level. A pop-up menu is presented that allows him to make the equivalent of a data base query. The result of this query is displayed in Figure 4. The user can interrogate the contacts by additional mouse interactions as shown in Figure 5.

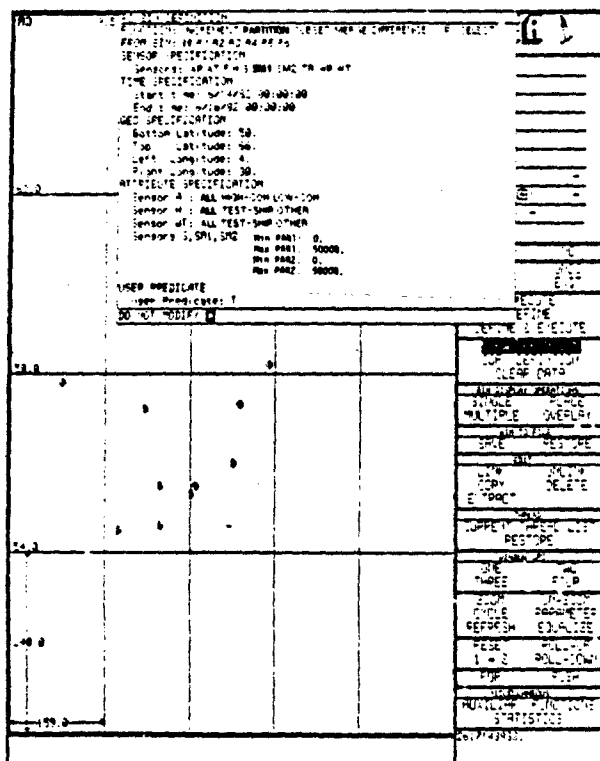




Figure 6 shows the menu of choices that the user has at the feasible track (FT) construction stage. In this stage, many feasible tracks are generated, each of which entail the feasible association of reports to tracks, along with scores measuring the quality of association. In addition to specifying the bins that he wishes to serve as data sources, the user can specify one of several correlation techniques. The methods currently available include:

- 1) A deterministic, algorithm based on kinematic feasibility and parametric gates.
- 2) A Bayesian scheme that considers both kinematic and attribute associations. The technique is based on the work of Bowman and Murphy [1].
- 3) An approximate Bayesian scheme which makes heuristic approximations to the complete Bayesian scorer.

In the first of these techniques, each report is associated to at most one track. In the latter two reports may be associated to more than one track, Kalmanfiltering is used, and scores are based on filter residuals. Also note that this menu contains several "control parameters", such as report parameter gate values, and a set of heuristic tree search parameter values, which control the combinatoric branching during the track generation process. [2] Figure 7 shows the output of this process for a Bayesian scoring scheme.

The final step is the generation of a consistent scene from the candidate tracks u at the FT level, which can be accomplished by uniquely associating reports to tracks. This is done at the FS level where a global optimization of the score is attempted. Typically, computational considerations preclude an exhaustive

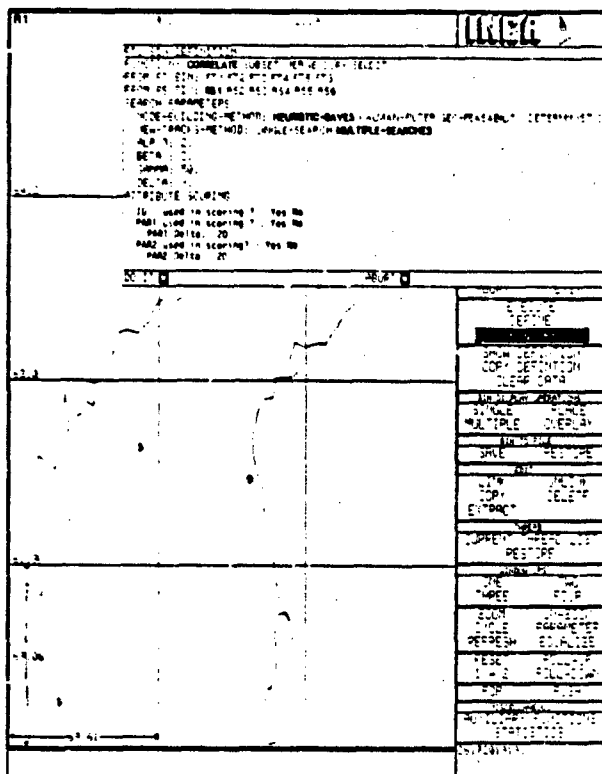


Figure 6. Choices for Feasible Track Construction

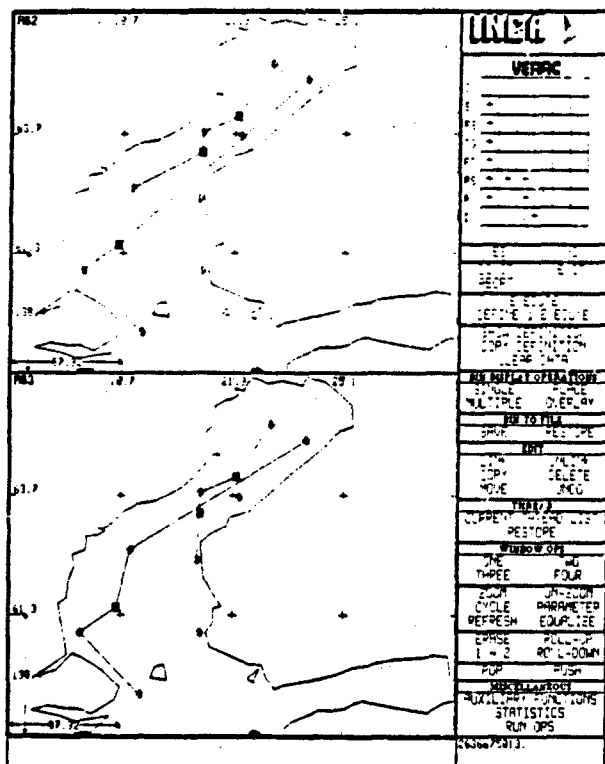


Figure 7. Candidate Feasible Scene

search, and we apply heuristic search methods as well. Figure 7 shows two feasible scenes scores that have been generated with a heuristic algorithm. The user would now select one or both of these to propagate with the next set of data.

The blackboard serves as an interface to the INCA data base where raw report, intermediate, and final products are stored. He may elect to process a particular set of data in several ways and the bins are used to store these results. There is no intrinsic hardware or software limitation on the number of scenes or number of intermediate results he can maintain. (The current graphical blackboard is limited to six scenes, but will be modified to display a structure that better maps to the tree that the user is building.) In addition to the algorithmic techniques available to the user, a large number of manual and graphic tools are available. Figure 8 shows one of the plotting options available for examining and creating data partitions via parametric value selection.

MOTIVATION AND METHODOLOGY FOR INTELLIGENT ASSISTANCE

The current INCA system allows the user to generate tactical pictures using a wide variety of tools and techniques. We are confident that this allows sophisticated users to understand better the impact of process improvements and to determine the best processing strategy for particular situations. Our assumption was that the analysts would apply very specific expertise relevant to each of the methods as they evaluated process performance and tuned the algorithms. Since the results of this tuning was to be made available to users, it was essential to capture this expertise. Thus, some other form of automated assistance in operating INCA was necessary. The conceptual framework for merging statistical decision theory and expert system technology has been explicated in [3].




Figure 8. Scattergram Facility

TRACK ASSOCIATION ALGORITHM

A.C. Bamford, E.M.L. Beale, D. East

SCICON Ltd., 49 Berners St., LONDON W1P 4AQ, ENGLAND

The Track Association algorithm combines a ship's underwater sonar and surface radar pictures in order to improve the quality of the underwater picture.

Introduction

The deployment of passive arrays has potentially greatly improved underwater detection performance, by making it possible to detect targets at much greater ranges than was previously possible. A consequence of this improved capability is that there tend to be many more targets within the array's detection range. This in turn complicates the task of processing observations, in order to derive an underwater picture consisting of tracks and their associated characteristics. The problem is made more acute by the fact that the observations carry limited information, namely bearing and frequency.

The formation of the underwater picture from array observations consists essentially of:

- assigning observations to tracks, that is deciding which observations come from the same target, and
- combining the information contained in the observations in order to form estimates of the tracks' characteristics.

The sonar operator can associate observations to form tracks provided that there has been no break in contact. However, he is unable to recognise when two or more tracks, which existed at different times, should be associated. This could arise, for example, after a manoeuvre by own ship, or as a result of sonar intermittency.

The Track Association algorithm aims to enhance a ship's underwater picture in the following ways:

- by recognising that several track segments originate from the same target, and
- by classifying track segments as ships, as a result of comparing track segments with the surface radar picture. If there are a lot of ships around then this is potentially a powerful way of improving the quality of the underwater picture, since the radar picture contains much more precise information.

The Algorithm

The algorithm combines a ship's underwater sonar and surface radar pictures into a single picture. We shall refer to tracks in the component pictures as segments and reserve the term tracks for tracks in the combined picture. A track therefore consists of one or more segments. The objective of the algorithm is to place segments from the same target in the same track, and segments from different targets in different tracks. The performance of the algorithm is the degree to which this can be achieved.

Parameterisation of Segments and Tracks

Segments and tracks are characterised in the algorithm by their estimated solutions and covariance matrices, which are parameterised in terms of Reciprocal Polar co-ordinates. These are defined as follows:

- ϕ_1 = the bearing of the target
- ϕ_2 = the inverse of the target's range
- ϕ_3 = the target's bearing rate
- ϕ_4 = the relative range rate of the target
- ϕ_5 = the frequency that would be received by a stationary observer at the origin in space and time (for sonar segments only).

All these quantities are measured from a stationary origin at or near the receiving sensor at a specific time.

The reason for using this parameterisation is that the covariance matrix provides more realistic confidence intervals for the estimates of sonar segments, in particular their estimated range. The parameterisation therefore provides a more accurate estimate of the penalty incurred (in terms of goodness of fit) when segments are combined to form a track. For a general discussion on Reciprocal Polar co-ordinates see [1].

Generation of Component Pictures

During the development of the algorithm we have generated the component underwater and surface pictures in the following way. First a scenario generator program

Initially, we utilized one of the expert system building tools emerging from AI laboratories to "capture" the analysts actions. The OPS5 system developed at Carnegie-Mellon University, was chosen. We have been utilizing OPS5 in this development activity for approximately 18 months.

Since many of the features of the INCA system are in continuous development, we did not process significant amounts of real data and the development of a sophisticated rule set was impossible. As a test of OPS5, we addressed the simpler problem of creating a system controller that is capable of running the INCA system, i.e., executing a sequence of actions and allocating machine resources. Eventually, this process will be expanded to include more complex decision tasks.

Up to this time, we have successfully built and tested a control system which uses simple resource allocation procedures (e.g., it always deletes the oldest contending entity) and can differentiate simple situations and select algorithms for feasible track generation.

REFERENCES

- [1] Bowman, C.L., C.L. Morefield and M.S. Murphy, "Multisensor Multitarget Recognition and Tracking," 13th Asilomar Conf. on Circuits, Systems and Computers, November 1979.
- [2] Nakagawa, S. and T. Sakai, "A Parallel Tree Search Method," Sixth International Joint Conference on Artificial Intelligence, Tokyo, 1979, pp. 628-632.
- [3] Bowman, C.L., L. Gross, H.J. Payne, "Data Fusion with Intelligent Assistance: INCA", 16th Asilomar Conference on Circuits, Systems, and Computers, November 1982.

generates the number and kinematics of targets

determines the positions of targets at specified times

derives raw observations of the targets consistent with the performance of ships' passive sonar and radars. (A sonar observation consists of a bearing and a frequency, while a radar observation consists of a bearing and a range).

assigns the observations to segments. Observations from the same target are assigned to the same segment provided that they are made by the same sensor and there has been no break of contact.

The estimated solution and covariance matrix of each segment is then calculated for each segment using a batch filter. This program minimises the sum of squares of residuals between fitted and observed values. Note that not all sonar bearing observations are processed since they cannot be assumed independent. Instead the first and last bearings plus bearings derived whenever targets change beam are processed. Note also that sonar segments may have two solutions resulting from the bearing ambiguity of sonar observations.

Assumptions

The initial implementation of the algorithm has made a number of assumptions, as follows:

- each target travels with constant velocity, consequently no attempt is made to detect manoeuvres, and
- each target emits a simple steady sonar frequency.

The assumptions were made to simplify the initial implementation, and will be relaxed in the next phase of development. It was considered preferable to solve the fundamental problems before being concerned with complicating factors. The scenario generator is consistent with the above assumptions.

Objective Function

The assignment of segments to tracks is determined by minimising an objective function. The objective function is defined by

$$C = S + \alpha T$$

where S is the residual sum of squares between fitted and observed values, taken over all tracks, T is the number of tracks and α is a constant. The value of α is chosen to balance two conflicting factors: the desirability of segments from the same target being associated and the undesirability of associating segments from different targets.

Assignment of Segments to Tracks

It would be too costly, in terms of computing time, to determine the absolute minimum of the objective function at specific times, so an approach which tries to improve on the existing assignment of segments to tracks is used.

The assignment of segments to tracks is reviewed periodically. At each review time a number of sets of segments are selected. Each set forms the basis of a review, and consists of segments which are assumed to have come from distinct targets. From our assumptions, the sonar or radar segments current at a given time satisfy this property. Sets are selected so that they include all segments current at the review time, all segments which started after the previous review time (in order to ensure that they are assigned to tracks at the earliest opportunity), and some (or all) historical segments.

The method of reviewing the assignment of the segments in a set to tracks is as follows:

- a. The characteristics of the reduced tracks, formed by removing the segments being reviewed from the current tracks, are determined.
- b. The increase in the objective function (or cost) resulting from assigning each segment to each reduced track is determined (i.e. the increase in residual sum of squares), as is the cost of not assigning a segment to any existing reduced track (i.e. α).
- c. Having worked out all the costs, the assignment of segments to tracks incurring minimum total cost is determined. This is done with the restriction that at most one segment can be assigned to any track (in order to be consistent with our definition of a set of segments).
- d. Finally the tracks are updated by assigning the segments to tracks in the manner determined in c. above. Segments which were not assigned to reduced tracks form single-segment tracks.

After completing the set reviews at each review time the tracks are examined to see whether any should be merged. Two tracks are merged if this results in a decrease in the objective function, and this happens if the increase in the residual sum of squares caused by the merge is less than the constant α .

Figure of Merit

We have defined a figure of merit (FOM) which takes values between -1 and +1, such that +1 is achieved when there is perfect association, and that the FOM decreases as the performance of the algorithm deteriorates.

Performance of Algorithm

The algorithm has been tested with several scenarios, each generating approximately 50 segments over a 12 hour period. The FOM for all these scenarios is +1 at every review time. It is worth noting that several segments (approximately 10%) consist of a single sonar observation (bearing and frequency) only.

In order to achieve this performance the method of computing the cost when assigning a segment to a reduced track, and the method of calculating the characteristics of a track from its component segments have undergone considerable refinement. An iterative process is now used which forms a trial solution, and then relinearises about this trial solution in order to form an improved solution.

Enhancements

During the next year we plan to relax the assumptions made during the initial implementation. We anticipate including the following features:

- Multiple sonar frequencies from a target
- Target manoeuvres
- Data from other platforms.

Acknowledgement

The development of the Track Association algorithm has been supported by the Admiralty Surface Weapons Establishment (ASWE) of the British Ministry of Defence.

Reference

- [1] The Use of Reciprocal Polar Co-ordinates in Passive tracking; A. C. Bamford, E.M.L. Beale, J. Lee, S. Patel; Proceedings of 6th MIT/ONR Workshop on Command and Control, 1983.

AD P002904

FURTHER RESULTS IN MULTIPLATFORM CORRELATION AND GRIDLOCK IN THE NAVAL BATTLEGROUP

M.A. KOVACICH

COMTEK Research, Inc., 100 Corporate Place, Suite B
Vallejo, California 94590

Command control and communication systems

ABSTRACT. Battlegroup ^{C3} requires that the force tracking picture developed by aggregate force sensors be accurate, nonredundant, and registered appropriately in each platform's coordinate frame. This report discusses further results in the development of systems, denoted Force Track Alignment systems, that satisfy these requirements. The evaluation methodology is discussed and performance data is presented that demonstrates the effectiveness of an FTA system that is composed of a Kalman bias estimation algorithm coupled to a sequential, track history dependent correlation process.

INTRODUCTION

Multiplatform correlation and gridlock, denoted in this report as Force Track Alignment (FTA), designates the distributed functional capability of the Naval battlegroup to align the local tracking picture of each platform with the cumulative tracking picture developed by the other platforms in the force. Alignment requires accurate, bias free coordinate conversions between platforms and reliable interplatform track pairings in order to present a one to one representation of the tactical environment to each platform in the coordinate frame of each platform. When viewed as part of the overall ^{C3} system of the battlegroup, FTA is a surveillance function in support of battlegroup operations. In current fleet operations, unfortunately, this support is limited and sometimes counterproductive due to the fragility and unreliability of the gridlock and correlation processes.

The Naval community has a number of ongoing projects to improve the FTA process in order to improve battlegroup coordination and to realize new engagement strategies, particularly intership fire control. The Gridlock Demonstration System (GDS), developed at Johns Hopkins University/Applied Physics Laboratory, is currently undergoing At-Sea testing. The FTA project, to be described in this report, is under NAVSEA sponsorship and is currently undergoing Landbased testing. In the near future, the JTIDS communication link, with its accurate relative navigation (RELNAV) capability, and NAVSTAR GPS, with its accurate geodetic navigation capability, will further contribute to the capability of the battlegroup to operate as a unit with surveillance assets on one platform directly supporting engagement requirements on other platforms.

The purpose of this report is to discuss some results obtained in developing FTA algorithms for deployment in a pre-JTIDS and pre-NAVSTAR GPS environment. The report is in two parts. The first part places FTA in the context of the Battlegroup ^{C3} system. The second part details the evaluation methodology and the performance results obtained in developing an FTA system.

FTA AS PART OF BATTLEGROUP ^{C3}

For the purpose of this discussion, Battlegroup ^{C3} will be viewed, in the sense of Morgan [1], as a nested, distributed collection of Perception - Assessment - Decision Taking - Execution (PADE) control processes, wherein:

Perception	is the process of sensing the environment and developing a representation of the state of the environment.
Assessment	is the process of relating perceptions to the objective of the control process.

Decision Taking	is the process of deciding on a course of action based on the assessment.
-----------------	---

Execution	is the process of carrying out the decision and interacting with the environment.
-----------	---

Generally, each PADE control process operates on a characteristic time scale over which the objective is constant and is coordinated with other PADE control processes that have overlapping characteristic time scales, objectives, or resource requirements.

Within this view of Battlegroup ^{C3}, FTA is essentially an element of the Perception process and directly supports the assessment and coordination activities of the battlegroup. FTA, through the process of intership track correlation, augmentation of remote data with local data on correlated tracks, link reporting responsibility logic, and intership bias removal strives to maintain a nonredundant force track file that contains the most complete and accurate track data generated by force sensors and registered in each ship's coordinate frame. Force coordination is supported by FTA since each platform perceives the same environment by means of the single, established force track file.

Improvements in FTA will lead to performance improvements in a number of PADE control process. To choose one, consider the timing relationships depicted in Figure 1 for the PADE control process that leads from initial radar detection by a surveillance radar to lock-on by the fire control radar. Suppose, also, that the initial radar detection occurs on the remote platform. (This situation is likely to occur in future battlegroup operation in which the remote platform carries the exceedingly accurate AN/SPY-1 radar and ownship carries a typical surveillance radar.)

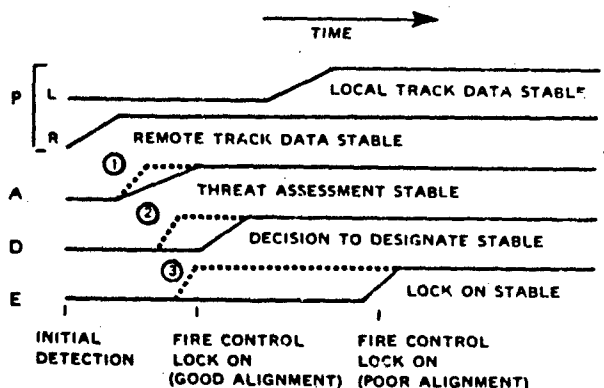


FIGURE 1. Timing Diagram for PADE Control Process leading from Radar Detection to Fire Control Radar Lock-On.

As the accuracy and reliability of the alignment process improves, a number of delays in this PADE control process can be removed:

- ① the time required for threat assessment can be reduced since the correlation decisions are more reliable leading to accurate threat assessments without as much corroborating evidence;
- ② therefore, the decision to lock-on can be made sooner, and
- ③ the fire control search algorithms need not wait for local track data since the intership biases are removed by the alignment process.

Effectively, when the alignment process is accurate and reliable, each platform can respond to the tactical environment as if each remote track is as accurate and reliable as a local track. When this capability is established in the Battlegroup, new engagement tactics (i.e., Decision Taking and Execution processes) become possible, particularly the capability to launch a missile on purely remote track data ("Launch on Remote" capability). These engagement tactics, generally denoted as intership fire control, will be particularly effective with the introduction of the AN/SPY-1 radar which is capable of highly accurate tracking. FTA will be a critical feature in realizing the full potential of this and other systems in Battlegroup C3.

Turning to a discussion of FTA itself, FTA designates the interrelated set of algorithms that automatically and continuously align the local track picture and the force track picture. As such, FTA correlates local and remote tracks received from all platforms on the data link, and estimates and corrects biases between each platform and the Gridlock Reference Unit (GRU) which establishes the standard coordinate system for the interchange of track data. The flow of data in an FTA system is shown in Figure 2 and discussed more fully in Kovacich [2].

Some underlying assumption made in the design of FTA for deployment in the 1984-1990 time frame are that each platform maintain a single set of bias estimates with the GRU via bias estimates with all platforms, and that FTA employ a single correlation algorithm to satisfy the dual requirements of generating mutual tracks with the GRU for girdlock purposes and generating virtual tracks with remotes from all platforms in order to maintain an accurate, nonredundant force track file. The technological environment is pre-JTIDS and NAVSTAR GPS and consists of current fleet radars and navigation systems and the current data link (Link-11).

The objective of the project is to find a mix of algorithms and algorithm thresholds such that FTA system performance is maximized. The objective is a formidable one given the variety of algorithm possibilities, the complex feedback interactions in FTA, and the large number of system thresholds.

To manage this problem, seven candidate FTA systems were designed which reflected various choices from a spectrum of possibilities (e.g., single pass vice multipass correlation constant g in vice Kalman bias estimation, uncoupled vice coupled systems, and FTA systems with and without open loop initialization algorithms). A Monte Carlo simulation program was written in order to stress the various FTA systems, to examine performance levels, and to optimize system thresholds. The approach to optimizing the system thresholds was to find the set of thresholds in which the FTA system exhibited acceptable correlation and bias estimation performance over the largest range of track densities in the presence of severe intership biases (10 nmi gridlock error, 4 deg azimuth bias, and 15 to 50 knots of intership drift). Standards for acceptable FTA performance were defined and specified in terms of allowed levels of dual designation (less than 5%), gridlock accuracy (less than .75 mmi), ini-

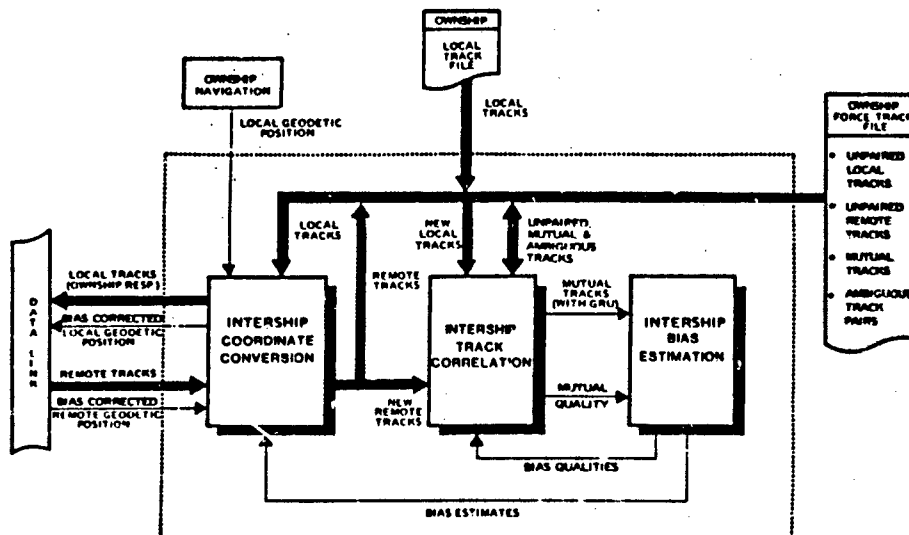


FIGURE 2. Flow of Track Data in the Force Track Alignment System. Only Closed Loop Stage is shown. Open Loop Stage generates Initial Bias Estimates and corrects Degraded Closed Loop Estimations.

ization time (less than 5 minutes), and stability across gation updates (less than 5% incorrect decorrelations). outcome of this phase of the evaluation program to maxi- each FTA system is presented in Figure 3. The core and ater load requirements of each system are plotted net the maximum track density for which the FTA system fied the acceptable performance standards. Each FTA em and all performance data are detailed in [4].




FIGURE 3. Maximum Track Density for which the FTA System demonstrated acceptable performance ($\times 10^{-5}$ tracks/NMI²).

he next stage in the evaluation program was to exercise ch FTA system with live radar data. A radar data tape was ovided by Applied Physics Lab/Johns Hopkins University and ntained time synchronous, smooth track data from two /SPS-39 radars, one at APL/JHU and the other at the Naval search Lab 34 miles away. The data tape is described in




FIGURE 4. Alignment Performance for Uncoupled FTA System that employs Single Pass Correlation and Constant Gain Bias Estimation (FTA 0).

FQ (dm)
(SCAN
AVERAGED
DISTANCE
BETWEEN
MUTUAL
TRACKS)




FIGURE 5. Alignment Performance of Loosely Coupled FTA System that employs Multi-Pass Correlation and Kalman Bias Estimation (FTA 8191.2).

detail in Miller [3]. Figures 4 and 5 present the performance of the best performing FTA system (FTA8191.2) and the poorest performing system (FTA0), respectively. The figures plot the scan averaged intertrack separation versus scan period. Note the comparative stability of FTA8191.2. The FTA0 suffered because of miscorrelations that drew off the bias estimates leading to further miscorrelations and bias errors. The miscorrelations and inability to later decorrelate the incorrect decorrelations were due mainly to the lack of internal coupling between the correlation and bias estimation process.

As a result of the simulation testing and system performance using the live radar tape, FTA8191.2, was recommended for implementation. The next section provides a description of this system.

DESCRIPTION OF RECOMMENDED FTA SYSTEM

The algorithm features of the recommended FTA system will be discussed in terms of Initialization, Correlation, Bias Estimation and Internal Coupling.

Initialization

The open loop initialization algorithm employs a two-track pattern matching algorithm that is used to provide an initial estimate of the intership biases and used in correcting bias estimates when the closed loop portion of FTA degrades through destructive feedback of miscorrelations to the bias filter. The algorithm searches for local-local track pairs whose spatial separation equals, within track noise, the separation between remote-remote track pairs. (Intertrack spatial separation between local tracks and between the corresponding remote tracks is translation and rotation invariant so is a valuable indicator that two local-remote track pairs are correlatable.) Additional correlation checks are made on the two local-remote pairs (e.g. identification, height and velocity tests) before finalizing the correlation. The set of correlations determined in this manner generate the initial set of innovations for the bias filter or the needed correction for the degraded bias filter.




FIGURE 6. Two-Track Pattern Matching. (L1,R1) and (L2,R2) are correlated only if DL and DR agree within Track Noise.

Correlation

The correlation algorithm carries out a sequential, two way (local-to-remote and remote-to-local) correlation process and makes correlation decisions on local-to-many remote ambiguities, remote-to-many local ambiguities, and many local-to-many remote ambiguities. (Currently, the algorithm makes correlation decisions on 2x2 matrix ambiguities only. Wait decisions are declared on higher order matrix ambiguities.) The correlation factor used to weight potential pairings and used in decorrelation processing for established pairings (mutuals) incorporates track positional histories and appears in Kovacich [5]. The correlation factor is recursively calculated and follows directly from Bayes rule:

$$CQ(t_i) = P(CORR, \bar{D}_L(t_i), \bar{D}_R(t_i))$$

$$= \frac{(L)(CQ(t_{i-1}))}{(L-1)(CQ(t_{i-1}))+1}$$

WHERE:

$$L = \frac{P[D_L(t_i), D_R(t_i) / CORR, \bar{D}_L(t_{i-1}), \bar{D}_R(t_{i-1})]}{P[D_L(t_i), D_R(t_i) / NO CORR, \bar{D}_L(t_{i-1}), \bar{D}_R(t_{i-1})]}$$

$$= (K) \{ \exp(-D^2 / (2\sigma^2)) \}$$

$$D = \text{SEPARATION BETWEEN } D_L(t_i), D_R(t_i)$$

$$\sigma^2 = \text{SUM OF LOCAL AND REMOTE TRACK VARIANCES}$$

$$K = \text{SCALING CONSTANT}$$

$$D_L(t_i), D_R(t_i) = \text{LOCAL, REMOTE POSITION VECTORS (SMOOTHED) AT TIME } t_i$$

$$CQ(0) = .5$$

Bias Estimation

The bias filter consists of two gains-limited Kalman filters: 1) the translation filter which estimates the latitude, longitude, latitude velocity, and longitude velocity biases and 2) the azimuth filter which estimates the intership azimuth bias. Both filters perform maneuver detection with gains reset in order to adapt to sudden shifts in intership biases.

The innovation sequence for the translation filter is the average lat/long separation vector between the local/remote tracks making up the mutuals held with the GRU. The innovation sequence for the azimuth filter is the average angular separation between line segments connecting local tracks and line segments connecting remote tracks. See Figure 7.




FIGURE 7. Definition of Innovation for Azimuth Filter.

The advantage of this definition of the innovation sequence for the azimuth filter is that it is translation invariant, thereby eliminating any coupling of the translation filter to the azimuth filter. Errors or transients in the translation filter have no effect on the azimuth bias estimate. On the other hand, the translation filter is coupled to the azimuth filter through the measurement covariance matrix, thereby treating the azimuth filter as a noise process.

Coupling

The bias filter and correlation process are coupled in two ways. First, the correlation process uses the bias filter covariance matrices to size correlation gates and to update the correlation probabilities. Second, the bias filter is coupled to the correlation process in that only mutuals whose correlation probability exceed a certain threshold are used in determining the innovations. This coupling aspect has proven to be very effective in generating a stable force track picture.

SUMMARY

This report discussed further results obtained in the design and evaluation of FTA systems. The relationship of FTA to Battlegroup C³ and the evaluation methodology were discussed and features of the recommended FTA system were presented. The recommended system is currently undergoing land-based testing in an operational Naval Command and Control program. At-Sea testing is to follow.

REFERENCES

- [1] Morgan, P.D. A Total System Approach to Command System Design. Proceedings of the 5th MIT/ONR Workshop on C³ Systems. LIDS-R-1267. December 1982.
- [2] Kovacich, M.A. A summary of Results in Multiplatform Correlation and Gridlock. Proceedings of the 5th MIT/ONR Workshop on C³ Systems. LIDS-R-1267. December 1982.
- [3] Miller, J.T. Gridlock Data Analysis - Initial NRL/APL Data. F2E-0-478 AFL/JHU. 22 June 1980.
- [4] Force Track Alignment System Selection Report. COMPTER Research, Inc. Report No. 1081E-2600-1. 1 June 1983.
- [5] Kovacich, M.A. An Aposterior Approach to the Multisensor Correlation of Dissimilar Sensors. Proceedings of the 4th MIT/ONR Workshop on Command and Control. LIDS-R-1159, Volume 1. October 1981.

INTERSHIP SENSOR ALIGNMENT USING UHF
TACTICAL DATA LINK RELATIVE NAVIGATION

R. H. Overton

Applied Technology Division
COMPTON Research, Inc.
596 Lynnhaven Parkway
Virginia Beach, Virginia 23452

Abstract

UHF tactical force data links can provide accurate relative navigation utilizing signal propagation delays, transformed into corrections to dead-reckoned relative position estimates by a decentralized extended kalman filter. The accuracy of this relative navigation can be exploited to achieve intership sensor alignment among the participating units. In certain common geometric situations, a simple range-only triangulation can be used. The results obtained by this method have accuracy comparable to more sophisticated filtering techniques, at a considerable reduction in computational time.

Introduction

A primary requirement for effective battle force coordination is the establishment of a united picture of the locations and actions of all militarily significant objects (tracks) in the tactical environment. In modern battle groups, with numerous units reporting surveillance information derived from diverse sensors, this requirement is frequently not met. Principal impediments to establishing a coordinated track picture are the inability of reporting units to locate themselves precisely in a consensual local navigation grid (gridlock); low data rates and unreliability of the communications links; misalignment among sensors on different platforms; and the inherent inaccuracies of the sensors themselves, including operator and tracker errors.

The use of frequency-agile, high data-rate UHF links, such as JTIDS, affords an opportunity to reduce the effects of poor gridlock and of slow, unreliable communications. Furthermore, the techniques involved are independent of the sensor errors experienced by participants. As a result, mathematical techniques can be used to correct for sensor misalignments, reducing sensor biases to the point of insignificance when compared to the random errors inherent in the sensors themselves.

The first section below outlines the method by which the Gridlock Problem is solved by UHF Relative Navigation. The second section proposes a highly accurate technique for correcting sensor misalignments among platforms. The third section contains an analysis of the remaining error in sensor alignment, due to all factors involved, including remaining gridlock errors and sensor system inaccuracy.

Relative Navigation

The tactical force data links using UHF signals can measure signal propagation delays for calculating ranges accurately enough to maintain good relative navigation (Rel Nav). An example of such a system is the Joint Tactical Information Distribution System (JTIDS).

A UHF link Rel Nav process employs two sets of data. The primary data are the outputs of on-board navigation sensors such as inertial or heading/speed systems, together with precise local clocks. The second set of data are "pseudorange," the propagation delays for Precise Participant Location Information (PPLI) messages, as measured by the recipients. These data, which include clock biases, are used to maintain a model of errors in the primary data. These errors are calculated in JTIDS by an 18-state Kalman filter maintained by each participant. Since each PPLI source reports his own position and time variances, these filters can be adaptive, and feedback is reduced by a convention that prohibits the use of PPLI messages with reported variances greater than the recipient's variances. To allow for both absolute and relative navigation to be passed on the link, an arbitrarily designated Relative (U, V) Grid is maintained by the participants. The geodetic positioning and orientation of this grid is updated by all units whenever any participant obtains a geodetic fix of his own position. The increased accuracy of his geodetic position report, as indicated in his PPLI messages, allows other participants to improve their estimates of their own geodetic position, either by direct propagation time-delay measurement or by relating the new geodetic information to the known position of the PPLI source relative to the recipient. This latter method allows the use of relayed (not line-of-sight) PPLI messages for improving absolute position estimates. Clock synchronization is enhanced by Round-Trip Timing (RTT) message exchanges (see Figures 1 and 2).





Figure 1. PPLI Messages



$$\frac{\Delta t_1 - \Delta t_2}{2} = \text{clock bias}$$

Figure 2. Round-Trip Timing

A critical factor in UHF link Rel Nav is the exclusion of data from sensors such as radar. This exclusion allows the separate computation of sensor bias by each tactical force unit, referenced to the common Rel Nav grid. The high levels of accuracy in Rel Nav, combined with elimination of large sensor azimuth biases, allow for much improved correlation/decorrelation criteria and a unified force track picture. Because of the independence of different units' calculations, no additional message traffic is required on the link.

Inter-Ship Sensor Alignment

A method for determining sensor azimuth bias is as follows: If U_0 holds another link participant, U_1 , on the sensor to be aligned, then the one-pass residual bias estimate using U_1 is simply

$$\Delta\theta = \tan^{-1}((x_0 y_1 - x_1 y_0)/(y_1 y_0 + x_1 x_0)) \quad (1)$$

where (x_0, y_0) is the position of U_1 according to U_0 sensors, and (x_1, y_1) is the position of U_1 according to his own report, expressed in an East-North (or U-V) aligned coordinate frame centered at U_0 (see Figure 3).




Figure 3. Using Rel Nav Position of Participating Unit

If U_0 does not hold another participant on the sensor to be aligned, U_0 can resort to a second method of estimating sensor bias, utilizing a mutual track held by U_0 and another link participant U_1 , as follows:

Let T be the object being tracked by both U_0 and U_1 , and reported by U_1 over the link. Let R_0 , R_1 be the distances from U_0 to T and from U_1 to T , respectively. Let D be the distance between U_0 and U_1 (see Figure 4).

Construct a rectilinear x' , y' grid at U_0 with positive x' -axis through U_1 . Then the position of T is calculated as:

$$x' = (R_0^2 - R_1^2 + D^2) / 2D \quad (2a)$$

$$y' = \pm \sqrt{R_0^2 - (x')^2} \quad (2b)$$

The choice of sign for y' is determined by the sign of ζ , the angle formed by the positive x' -axis and the ray from U_0 to T .




Figure 4. Utilizing Mutual Track

Error Analysis

First, consider the general error model for any one-pass azimuth bias estimate. The error ϵ_θ is, in general, a function of five random variables:

- X_0 = noise in measuring U_0 azimuth (radians)
- X_1 = noise in relative position estimates by U_0 , U_1 (feet)
- X_2 = noise and bias in U_1 azimuth (radians)
- X_3 = noise and bias in U_0 range (feet)
- X_4 = noise and bias in U_1 range (feet).

These five random variables are assumed to be independent. X_0 and X_1 have mean 0.

For each residual bias estimate using equation (1), ϵ_θ can be approximated linearly by

$$\epsilon_\theta = X_0 + X_1 / D \quad (3)$$

If equations (2a and 2b) are employed, let γ be the angle between bearings to T from U_0 and U_1 . Then

$$c_0 = X_0 + X_1 / D + X_3 \cot \gamma / R_0 + X_4 \csc \gamma / R_0 + S(X_0) \quad (4)$$

where $S(X_0)$ is the error in choice of sign for y' .

From equations (3) and (4), it follows that if equations (1) or (2a and 2b) are used to calculate residual bias, then:

1. c_0 is independent of X_2 ;
2. If equation (1) is used, the estimate is the optional measurement, since any measurement includes errors X_0 and X_1/D ;
3. The coefficients for X_3 and X_4 are bounded by $1/R_0 \sin \gamma$.

For many sensors (e.g., radars), range errors are small compared to the ranges at which they operate. In these cases, this bound implies that, with suitable geometry constraints, the X_3 and X_4 terms of (4) are no greater than X_0 and frequently are much less.

4. The error $S(X_0)$, resulting from the choice of the incorrect sign in Equation 2b, can be handled by observing that $S(X_0) = 0$ or $S(X_0) = -2y'/R_0$, depending on the choice of sign for y' . Therefore, for any fixed value $y_0 > 0$, one can require that $|y'| > y_0$ in Equation 2b; otherwise, the track T will be discarded from further processing. If $S(X_0) \neq 0$, then $|S(X_0)| > 2|y'|/R_0 > 2y_0/R_0$. Therefore, if one rejects any residual bias estimate, $\Delta\theta$, greater than, say, y_0/R_0 , then for any accepted $\Delta\theta$, the probability of $S(X_0) \neq 0$ is less than the probability that the total sensor azimuth error ($X_0 + \text{bias}$) exceeds y_0/R_0 . Thus, the probability of $S(X) \neq 0$ can be made arbitrarily small by choosing y_0 to be sufficiently large; the choice depends upon the distribution of the total errors experienced by sensors of the type being aligned.

The proposed bias estimates are thus expected to give errors on the order of $X_0 + X_1/D$ under minimal geometric constraints. If the variance of this expression is given by σ^2 , then the least-squares approximation using n observations will have an error variance of approximately σ^2/n . When implemented in varied simulated environments, the methods described above exhibit errors within these expected values, with considerably lower processing than required by a least-squares filter which models the (x, y) position of U_0 , U_1 , and T for each track.

THE USE OF RECIPROCAL POLAR CO-ORDINATES IN PASSIVE TRACKING

A.C. Bamford, F.M.L. Beale, J. Lee, S. Patel

SCICON Ltd, 49 Berners St., LONDON W1P 4AQ, ENGLAND

The use of Reciprocal Polar co-ordinates to parameterise tracks is discussed. Their use makes the problem of determining the characteristics of tracks from bearings only data much less non-linear than the use of the more conventional Cartesian co-ordinates.

Introduction

The task of tracking a moving target from sensor data can be formulated by assuming that the target moves with constant velocity. Once this basic problem has been solved, we may extend it by assuming that the target manoeuvres from time to time, that is to say it changes its speed or course. For this basic problem it is natural to represent the position of the target in terms of parameters ξ_1, ξ_2, ξ_3 and ξ_4 such that the Cartesian co-ordinates of the target at time t are:

$$(X_0 + \xi_1 + \xi_3(t - t_0), Y_0 + \xi_2 + \xi_4(t - t_0))$$

where (X_0, Y_0) denote some convenient origin, and t_0 is some convenient time. We can then compute fitted values for the observations of bearing, range or rate of change of range, and hence find parameter values to minimise some weighted sum of the squares of the discrepancies between the measured and fitted observations.

However, if most of the data consist of bearings from a single slowly moving platform, then the model is highly non-linear in the sense that, over the whole range of plausible target tracks, the fitted bearings are highly non-linear functions of the parameters. This nonlinearity can largely be removed by using a different parameterisation. We shall refer to the parameterisation we chose to reduce the nonlinearity, as Reciprocal Polar co-ordinates.

In the course of this paper, we outline our reasons for choosing this parameterisation in the first place, and then proceed to describe the work we carried out to confirm our ideas. In the main this consisted of calculating the non-linearities for a particular scenario with respect to various parameterisations.

1936 Formula

Beale [1] suggested and Bates and Watts [2] confirmed that many nonlinear regression problems, of which the Bearings Only problem discussed here is an example, can be made much less nonlinear by making a suitable nonlinear transformation of parameters. Tenney et al [3] discuss a related problem, and propose a transformation that greatly improves on Cartesian co-ordinates.

An appropriate transformation for the Bearings Only problem is suggested by considering the standard 1936 Formula for estimating a target's range. If the observed bearing rate is B_1 when own ship moves with a velocity U_1 across the bearing and is B_2 when own ship moves with a velocity U_2 across the bearing, then if the target is at range R and moves with a velocity component V at right angles to the bearing, we have

$$B_1 = \frac{V - U_1}{R} \quad \text{and} \quad B_2 = \frac{V - U_2}{R}$$

or, in other words

$$\frac{1}{R} = \frac{B_1 - B_2}{U_2 - U_1} \quad \text{and} \quad V = \frac{U_2 B_1 - U_1 B_2}{U_2 - U_1}$$

This means that the reciprocal of the range and the bearing rate from a stationary point near own ship are linear functions of the data B_1 and B_2 .

Proposed Parameterisation

The above analysis suggests that the parameters ξ_1, ξ_2, ξ_3 and ξ_4 should be replaced by parameters ϕ_1, ϕ_2, ϕ_3 and ϕ_4 representing the bearing and the reciprocal of the range of the target at time t_0 from a point at or near own ship's position at this time, the bearing rate, and some parameter related to the range rate. Our first thought was that ϕ_4 should be the rate of change of the reciprocal of the range. But further work, described later, suggests that ϕ_4 should be relative range rate, that is the velocity of the target along the bearing line divided by the range. It is natural to call these Reciprocal Polar co-ordinates. To use them we need formulae for the ξ_i as functions of the ϕ_j and for elements of the Jacobian matrix $(\partial \xi_i / \partial \phi_j)$. Given formulae for linear approximations to the fitted values of any observation in terms of the ξ_i , we can then derive corresponding linear approximations in terms of the ϕ_j .

Batch Filter

We have implemented a "batch filter", based on Reciprocal Polar co-ordinates. It is an iterative non-linear least squares procedure for estimating target tracks from passive bearing data, and calculates the sums of squares of residuals in terms of Reciprocal Polar co-ordinates. The program currently only considers non-maneuvring targets. The batch filter is used extensively by the Track Association Algorithm [4].

Special precautions taken in the implementation of the batch filter are as follows.

- (1) There is a need to impose both a lower and an upper bound on ϕ_2 (the reciprocal of the target's range). The lower bound is to ensure that the target's range remains realistic while the upper bound avoids the singularity at the origin. These bounds can be imposed by pivoting last on this parameter and modifying the data before the pivot so that it will take the desired boundary value. The method of pivoting is described by Stiefel [5].
- (2) It is necessary to have a safeguard against non monotonic behaviour in the residual sum of squares, and can be achieved using a simple bisection formula.

Benefits

Reciprocal Polar co-ordinates are superficially more complicated than their Cartesian counterparts, but we suggest they offer the following advantages.

- (1) The task of computing true least-squares estimates of the target parameters is simplified, since we need fewer iterations and can dispense with the usual Levenberg-Marquardt modification to the Gauss-Newton method. In general only 3 or 4 iterations are required to get close to the least squares solution. (This solution is not, of course, dependent on the parameterisation used).
- (2) The parameterisation allows the computation of realistic confidence intervals for the target parameters, and in particular the current target range. This confidence interval may be unsymmetrical, since we will often know that the target cannot be very close but may be much further away than our best estimate. This information is particularly relevant when one has to decide whether 2 tracks can reasonably be attributed to the same target.
- (3) Finally, a nearly linear parameterisation allows some use of sequential (that is non iterative) estimation procedures, and in particular the use of manoeuvre detectors based on cumulative sums of the innovations recommended by Brown et al [6] and Patel [7].

We have verified these theoretical advantages to the extent of having developed the batch filter, which we outlined earlier, and a program to compute the non-linearity for various parameterisations of a typical bearings only scenario. In the next section we discuss this computation in detail.

Calculation of Non-linearity

The possibility that a non-linear transformation of parameters can make a non-linear regression model much less non-linear was mentioned earlier. This suggests that non-linearity can be split into two types: Beale [1] called these Intrinsic Non-linearity and Parameter Effects Non-linearity. Beale defined formulae for non-linearity such that N_I (the intrinsic non-linearity) is the minimum measurement of non-linearity taken over all parameterisations of the model, while N_R (the parameter effects non-linearity) is the total non-linearity for a given parameterisation minus N_I .

We calculated N_I and N_R for a given scenario, firstly for Cartesian co-ordinates, and secondly for the family of co-ordinate systems satisfying

$$\begin{aligned}\phi_1 &= B \\ \phi_2 &= R^{\alpha_1} \\ \phi_3 &= B \\ \phi_4 &= R^{\alpha_2} R\end{aligned}$$

where B and R are the target's bearing and range, and α_1 and α_2 are constants. Note that we have Reciprocal Polar co-ordinates when α_1 and $\alpha_2 = -1$.

The scenario is illustrated in Figure 1.

CALCULATION OF NON-LINEARITIES - SCENARIO




Figure 1

The ship travels on a dog-leg course at 10 knots making a bearing observation every 10 minutes. Observation errors are normally distributed with a standard deviation of 2°. Altogether the ship makes 9 observations. The target submarine is initially 50nm from the ship and is travelling at a constant 20 knots.

The intrinsic non-linearity for this scenario is 0.00191, while the removable non-linearity for Cartesian co-ordinates is 0.547. Figures 2,3 and 4 show the value of the removable linearity for the family of co-ordinate systems for various values of α_1 and α_2 . These figures indicate that the parameterisation with $\alpha_1 = \alpha_2 = -1$ (i.e. Reciprocal Polars) gives a minimum of the removable non-linearity for the scenario. Furthermore the total non-linearity using Reciprocal Polar co-ordinates is more than a factor of 12 less than the corresponding figure for Cartesian co-ordinates. These results justify our original reasons for defining Reciprocal Polar co-ordinates.




Figure 2




Figure 4




Figure 3

Further Work

The work described in this paper has investigated the non-linearity of various parameterisations when the observed data consist of bearings only. However, the bearing observations are often made in conjunction with frequencies which can provide information on a target's closing speed.

We intend in the near future to calculate the non-linearities of various parameterisations when frequency information is available. Such parameterisations do of course require a fifth frequency-related parameter. It is perhaps worth noting that our batch filter based on Reciprocal Polar co-ordinates has no problems converging when processing data including frequencies.

We shall also be extending the batch filter so that it can track manoeuvring targets by making use of innovations.

Summary

We have proposed the use of Reciprocal Polar co-ordinates as a means of reducing the non-linearity of the Bearings Only tracking problem. The co-ordinate system is suggested from considering the 1936 Formula, and was justified more formally by calculating the intrinsic and parameter effects measures of non-linearity for a given scenario using various parameterisations. Practically, we have implemented a batch filter, based on Reciprocal Polars, which minimises the sums of squares of residuals. The filter works well and converges quickly.

References

- [1] E.M.L. Beale (1960) "Confidence regions in non-linear estimation" J. Roy. Statistical Soc. (B) 22 pp 41-88.
- [2] D. M. Bates and D. G. Watts (1980) "Relative Curvature Measures of Non-linearity" J. Roy. Statistical Soc. (B) 42 pp 1-25.
- [3] R. R. Tenney, R. S. Herbert and N. R. Sandell (1977) "A Tracking Filter for Manoeuvring Sources" IEEE Transactions on Automatic Control. April 1977 pp 246-251.
- [4] A. C. Bamford, E.M.L. Beale and D. East (1983) "Track Association Algorithm" Proceedings of 6th MIT/ONR Workshop on Command and Control, 1983.
- [5] E. L. Stiefel (1963) "An Introduction to Numerical Mathematics" Academic Press.
- [6] R. L. Brown, J. Durbin and J. M. Evans (1975) "Techniques for testing the constancy of regression relationships" J. Roy. Statistical Soc. (B) 37 pp 149-192.
- [7] S. B. Patel (1982) "Real Time Passive Tracking For Tactical Applications" Journal of Naval Science Vol 8, No 3 1982 pp 192-201 (Restricted).

STOCHASTIC CONTROL OF A PASSIVE SEARCH FOR OCEAN VEHICLES

LCdr. G. Z. Wilhoit
U.S. Navy

Prof. R. R. Tenney
MIT/LIDS
Rm. 35-213

Abstract

The problem of selecting a maneuver policy for single-platform bearings only search and localization in the ocean is viewed as a stochastic optimal control problem. A solution is developed based on a realistic environment model and extended Kalman filter tracker. Simulation results highlight the dominance of engagement geometry (e.g. direct-path vs. convergence zone) over maneuver policy with respect to total performance.

I. Introduction

1.1 The Bearings-Only Problem

Due to the complex nature of sound propagation in the sea, a passive acoustic search involving ocean vehicles inherently provides a minimum of information of a very noisy nature to the searcher. Currently the research emphasis seems to be toward developing methods to utilize the acoustic information that is available in a more efficient manner. For instance, acoustic doppler information can often be gleaned from the signal spectrum of a target, providing range rate information on the contact. While this knowledge of range rate can be a tremendous aid in specifying the target's state (range, bearing, course and speed), it is operationally easy for an adversary to introduce sufficient randomness in the frequency spectrum emitted by his vessel to make his acoustic signal unreliable for this purpose. Another area of interest has been to increase the acoustic aperture of the hydrophone arrays on board present searching units. By measuring the time differences between the acoustic signals incident to each hydrophone combined with the array's bearing resolution, passive target localization is possible. However, with current capabilities, such an increased acoustic aperture is too large for feasible installation on searching vessels. Thus, the acoustic information available on each source is often limited to frequency, sound pressure level (SPL) and bearing. Frequency information can be extremely valuable for use in target classification, but for tracking and especially detection functions, it has many limitations. Due to the propagation vagaries of sound in the sea, SPL information may also be quite misleading. Consequently, bearing information, although potentially noisy, provides the best single means of target detection and tracking with current passive acoustic sensors in the ocean environment.

1.2 Approach

The approach to the passive acoustic search problem followed in this paper consists of three major parts. First, a realistic model of the ocean environment including the searcher and target vehicles will be developed in order to simulate the stochastic and dynamic behavior in the "real world". The model will include the behavior inherent in its four major components: the target, the medium, the searching vehicle and the searcher's sonar equipment. Second, an estimator will be developed which will process information provided by the model as an aid in making various search decisions. This estimator will be

exercised principally in a localizing search of a given area to provide the best estimate possible of a target's course, speed and range at any instant in time. Finally, based on all of the information provided from both the model and the estimator, a strategy will contain all of the necessary control logic to conduct the search.

With these three major components connected properly a series of strategies which seem reasonable can and will be evaluated in an open-loop manner by means of computer simulation. By testing these strategies representing various scenarios of interest, the advantages and disadvantages of each can be determined for different types of vehicle encounters. Subsequently, the desirable portions of each strategy can be implemented with the required decision logic to conduct closed-loop simulations for final evaluation.

II. The Model

The ocean environment is an extremely inhomogeneous and unpredictable medium in which to conduct a search. The acoustic interactions between the propulsion machinery and sonar equipment on board target and search vehicles in these surroundings make the passive search and tracking problems extremely difficult to describe accurately yet succinctly. Although more difficult than many systems, the ocean vehicle search problem can be approached in a conceptual way as a complicated control problem. As in the control of any system, it is necessary to establish a model which is a trade-off between two conflicting desires. The first and most obvious desire is to establish a model which is detailed enough to adequately represent the actual system to be controlled. Secondly, an equally necessary characteristic of a useful model is that it is simple enough to "solve" for a practical control strategy. As the complexity of a system increases, finding an acceptable model which satisfies both of these desires becomes more and more difficult.

2.1 Environment

An examination of the passive sonar equation can provide a model which reasonably approximates most interesting areas of the ocean. In general, the most important components of the transmission loss will be spreading, (spherical and cylindrical), absorption, leakage and convergence zone effects. A reasonable approximation for most areas is the transmission loss expression for the mixed-layer case corrected for convergence gain effects at the appropriate convergence ranges.

The average values of the various sonar parameters in the sonar equation can be used to compute a probability of detection function to be used in the actual model implementation. This conversion from average values to some type of probability distribution or statistic is necessary to model the intermittent directions made at the sonar system output for time-varying SNR's near zero. From Knight's [3] description of current digital signal processing technology used in sonar systems, a graph of probability of

detection versus SNR can be derived (see Section IV).

2.2 Vehicle Dynamics

In this model, target courses and speeds can be modeled in a variety of ways. While in this application they were deterministic functions of time, they could be determined by any other desired function of deterministic or random variable. In the present case, the dynamics of each vehicle were computed by

$$X(i+1) = X(i) + L \cdot VS \cdot \sin(VC)$$

$$Y(i+1) = Y(i) + DT \cdot VS \cdot \cos(VC)$$

where X and Y = the cartesian coordinates of the vehicle

DT = the time difference between the i and $i+1$ observations

VS = vehicle speed

VC = vehicle course.

2.3 Observations

The actual bearings between searcher and target were then computed by

$$AB = \arctan((YT-YS)/(XT-XS)),$$

where XT and YT = the cartesian coordinates of the target

XS and YS = the cartesian coordinates of the searcher.

These actual bearings do not represent bearings that realistically would be observed in a noisy environment. In order to model these observed bearings a zero-mean, Gaussian noise component with a standard deviation of .1 degrees was added to each actual bearing. The statistics of this observation noise can be relatively accurately predicted from the detection equipment used; it implies an observation noise variance of $R = .01$.

2.4 Probability of Detection Versus Range

Besides the large number of measurements that must be made, a limitation of another kind is produced by the nature of the medium in which sonars operate. The sea is a moving medium containing inhomogeneities of various kinds, together with irregular boundaries, one of which is in motion. Multipath propagation is the rule. As a result, many of the sonar parameters fluctuate irregularly with time, while other change because of unknown changes in the equipment and the platform on which it is mounted. Because of these fluctuations, a "solution" of the sonar equation is no more than a best-guess time average of which is to be expected in a basically stochastic problem. Precise calculations, to tenths of decibels, are futile. A predicted sonar detection range is an average quantity about which the observed values of range are likely to congregate. Thus, the model supplies the probability of detection versus range for the target of interest to be used on-line for target state estimation and strategy formulation purposes. The calculation of a specific probability of detection versus range curve is included in Section IV for use in various open-loop simulations.

III. The Estimator

There are many estimator designs in the literature which have been applied to various search and tracking problems. For the passive ocean search problem there are several factors which will influence the selection of the estimator. For a bearings-only search and track scenario the dynamics of vehicle motion are represented by nonlinear equations. In this case involving imperfect information, the implementation of such an estimator is by no means easy.

3.1 Bearings-Only Tracking

The bearings-only case of target tracking introduces some unusual problems. The many sources of noise which may affect received bearing information in the ocean environment have already been discussed. Most of the other problems associated with bearing-only tracking involve the specific geometries which are involved. Convergence of filters used for bearings-only tracking depends on many factors. Target trajectories involving long ranges and small velocities may make accurate solutions take longer to converge since the change in bearing with time is quite small [5].

Another condition of geometry which has been shown to have a very profound effect on tracking performance in the bearings-only case is the choice of coordinate system for estimator implementation [6]. Since the final filter configuration for any specific problem will ultimately depend upon which reference frame is employed during problem formulation, it is not surprising to find that Cartesian coordinates are used extensively to formulate target motion analysis (TMA) estimation problems in the context of the EKF. This reference frame permits a simple linear representation of the state dynamics: all system nonlinearities are embedded in a single scalar measurement equation [7]. Such a modeling structure is especially appealing for practical application because it minimizes filter computational requirements.

Many non-Cartesian filters have been used which possess significantly different, and perhaps better, performance characteristics than their Cartesian counterparts [8]. In fact, there is new theoretical and experimental evidence that seems to indicate that Cartesian filter implementations may be unstable for single sensor bearings-only TMA [4]. Specifically, these results indicate that unique interaction and feedback of estimation errors within this filter render it highly susceptible to premature solution divergence and covariance collapse.

3.3 Estimator Architecture

The estimator chosen to conduct bearings-only TMA in this application consists of a three-dimensional array of Extended Kalman Filters with hypothesis testing occurring along each of the three dimensions. By assuming that the bearing noise is zero mean and Gaussian, the EKF may be used to provide estimates of range, course, and speed errors given bearing observations in a large fraction of possible cases. The problem is that in the bearings-only case, initial course, speed and range information is not readily available and in order for an EKF to be effective, it must be initialized relatively close to the actual solution. In order to overcome this problem we employ the technique of hypothesis testing, hypothesizing courses, speeds, and initial ranges. Two of the dimensions of the array of EKF's will be based on hypothesized discrete courses and speeds which span the set of possible targets. The filters will then generate estimates of the differences in course and speed between the actual and hypothesized values. The third dimension of the estimator will consist of three initial range estimates, two based on direct path and one or two based on convergence zone range of detection predictions. Thus by observing only bearings with noise, this array of EKF's will estimate ranges and bearings to the target and differences between estimated and hypothesized courses and speeds of the target. By hypothesis testing based on the covariance matrices computed during the tracking evolution, some filters will be turned off and the "most likely" filters to be "correct" will be left running.

3.3 The Extended Kalman Filter

The Extended Kalman Filter Equations are as follows
[2]. The state vector used in both mean prediction and mean update steps is

$$\underline{X} = \begin{bmatrix} R \\ B \\ DC \\ DS \end{bmatrix}$$

where R = target range from searcher
B = target true bearing from searcher
DC = delta course from hypothesized target course
DS = delta speed from hypothesized target speed

The reason for picking a non-Cartesian state vector have previously been explained. Finding the optimal non-Cartesian coordinate system was not the goal of this effort. This coordinate system is used since it uses the same four parameters (range, bearing, course, and speed), that are used operationally and thus are the most logical parameters to extend directly into strategy formulation without needing prior transformation.

Since the state dynamic equations are nonlinear, a linear approximation to the \underline{A} matrix used for covariance prediction must be used. A first order linear approximation was used such that the rows of \underline{A} are the partial derivatives of the predicted state with respect to the previous state vector.

The prediction includes the normal addition of uncertainty \underline{Q} to the covariance after transformation to account for the dynamics. An additional term, \underline{Q} , is necessary for the nonlinear covariance prediction to account for higher order terms in the Taylor Series expansion that were neglected in obtaining the \underline{A} matrix. As used in this application, \underline{Q} and \underline{Q} were assumed constant and an approximation to their sum was obtained empirically. In order to find this sum a series of sample runs was conducted with the searcher in various opening and closing scenarios representing the full range of searcher-target aspects. A circumnavigation of the moving target vehicle by the searcher was primarily used since it exhibited a wide variety of angles of encounter. First, an estimate of each diagonal term in the \underline{Q} matrices was made that was about two orders of magnitude below the initial errors to be expected in practice and used in the simulations. With these values serving as an initial baseline from which to start, runs were made while varying these values one at a time three orders of magnitude above and below these original values.

From these initial tests, a first approximation of the "best" value for each of the four diagonal terms was determined by comparing the diagonal elements of the updated covariance matrices generated with the actual errors on each state squared. The criterion was to pick \underline{Q} values which caused the computer filter covariance matrix to best match the magnitudes of the squares of the state errors. These new values of \underline{Q} served as a more precise baseline from which each of the values of \underline{Q} were varied again. By the fourth baseline of \underline{Q} values, the \underline{Q} matrix sum was within an order of magnitude of the "best" value on each diagonal term and the covariance terms were within a factor of three of the squared state error terms. The final value obtained from this iterative optimization procedure was

$$\underline{Q} + \underline{Q} = \begin{bmatrix} .0001 & & & \\ & .00001 & & \\ & & .0001 & \\ & & & .00001 \end{bmatrix}$$

where a typical resulting filter covariance matrix and squared state error matrix were

$$\underline{P} = \begin{bmatrix} 1.818 & & & \\ & .01117 & & \\ & & .3875 & \\ & & & .5918 \end{bmatrix}$$

and

$$\underline{dx}^2 = \begin{bmatrix} .6478 & & & \\ & .00376 & & \\ & & .1633 & \\ & & & .3173 \end{bmatrix}$$

Again, the non-diagonal elements of the \underline{Q} matrix were assumed to be small and were neglected.

3.4 Hypothesis Testing

All of the computations thus far have assumed a prior knowledge of the target's course and speed. This information is provided by hypothesizing various combinations of course and speed, one for each Extended Kalman Filter. Initially, twelve courses and twelve speeds were hypothesized to test the filter. The hypothesized courses started at 000 degrees true with increments every 30 degrees and the speeds were at increments of 4 knots starting at 0 knots. As the final implementation including the sum of the \underline{Q} matrix terms became clear, a review of the filter's ability to distinguish errors in course and speed was conducted to determine if this initial discretization of hypothesized courses and speeds was valid.

A criterion of having the hypothesized values spaced no more than four times as far apart as the average filter estimate errors seemed to ensure adequate estimator filter overlap. From the test runs conducted with the appropriate \underline{Q} matrix sum terms, average absolute errors in course and speed of 7.5 degrees and 1.5 knots were indicated. Therefore the final hypothesized courses were taken every 30 degrees in a circle starting at 000 degrees true for a total of 12 different courses. The hypothesized speeds were taken at increments of 6 knots from 0 to 42 knots for a total of 8 different speeds.

The final estimator consisted of three banks of 96 Extended Kalman Filters with all of the expected courses and speeds discretized in each bank. Two banks of filters were initialized at two direct path ranges of detection. Two initial direct path ranges were chosen since by adding the range estimate error determined to be acceptable for the EKF performance during trial runs to these two initial ranges, the total encompassed range band approximated the expected errors on the range prediction provided by the sonar equations. The third filter was initialized at the predicted range of detection in the first convergence zone. Since the first convergence zone is the narrowest convergence zone, only one initial range was required to match the EKF's range error performance to its expected width.

The only remaining requirement to implement this hypothesis testing scheme is to pick some statistic generated by each of the 288 filters on which a decision, "correct filter" or "incorrect filter" could be based. The statistic used was

$$\chi^2 = \text{RES}^2 / (R + P(2,2))$$

where RES = the residual as computed during the mean update

R = the bearing observation noise variance

P(2,2) = the bearing covariance term from the covariance prediction.

For each subfilter, this computation would result in a χ^2 (chi-square) density function with one degree of freedom if the problem were truly linear and Gaussian [9]. For the correct filter, the expected value of this random variable is unity, indicating that the

actual bearing errors are of the same magnitude as the predicted bearing errors plus the observation noise. This statistic should have a higher value for all other incorrectly matched filters. In order to smooth this performance test and look at the recent filter history, the actual statistic implemented was

$$x(t)^2 = \alpha * RES^2 / (R + P(2,2)) \\ + (1 - \alpha) * x(t-1)^2$$

where by setting $\alpha = 1$, the most recent ten minutes of performance were evaluated. The logic used to turn off diverging filters was set in the range $X^2 [1,5]$. This range of value was chosen to obtain the "best" 10-30 filter solutions by the end of the run.

IV. The Strategy

4.1 Simulation Options

In comparison with decisions already made concerning model and estimator implementation, picking a limited number of meaningful simulation scenarios was much more difficult. The clear desire in this case was to select realistic strategy options which would provide the most "insight per run" and still keep them simple enough so as to minimize any coupling which could cloud their relative benefits.

One of the most important decisions a searcher must make after an initial detection is how to maneuver so as to acquire as much information about the contact as possible in the shortest period of time. In the bearings-only case, it is clearly important for the searcher to maneuver, influencing the bearing information received in such a manner that the set of feasible target states is reduced as much as possible. Often such maneuvers are limited to course rather than speed changes since the searcher would like to run at a speed which due to self-noise provides him the highest acoustic advantage. Thus, a simplified version of this strategy decision would be at what angles and at what rate should course changes occur for optimal target tracking.

A second and slightly more obscure factor of interest is the initial aspect of encounter between the target and searcher. At first glance, this factor appears to be a poor simulation variable since under most conditions, the searcher has very little control over the manner in which he first confronts the target. However, the initial geometry of target detection is important since by studying tracking performance under these different initial conditions, the more advantageous target-searcher aspects can be determined. By using this information, the searcher can improve his tracking performance by maneuvering into a more favorable target-searcher aspect than he first encountered.

The factors already discussed are important primarily due to their geometries. Another factor which is of particular importance due to the complex acoustic environment already discussed is the propagation path by which the signal is received. Based only on bearing information, the difference in the initial range of detection between a direct path and a convergence zone contact could be a factor of 4 times or more. We would expect the most desirable strategies to be quite different between the different path and convergence zone propagation cases with range differences this large. Therefore, the propagation path is also included as one of our simulation variables.

4.2 The Simulation

The simulation consisted of a maneuvering searching unit and a constant velocity target. The simulation was run using three different target-searcher aspects:

1. initial detection on target's beam
2. initial detection 45 degrees off the target's bow
3. initial detection 45 degrees abaft the target's beam

This discretization of relative aspects was chosen so as to cover a major portion of the target's azimuth while holding the increase in the required number of runs to a factor of 3.

The searcher pointed the initial contact bearing and commenced a zig-zag search on that base course. Three zig angles were used from the base course.

1. 30 degrees
2. 60 degrees
3. 90 degrees

In addition, three different zig times were used during the simulation.

1. maneuver every 5 minutes
2. maneuver every 10 minutes
3. maneuver every 15 minutes

These maneuver angles and frequencies were chosen since they seem to bound the most interesting range of values. Each of these 27 runs was made under two different initial range conditions. One run was made with the initial detection made at a range predicted by direct path acoustic propagation. The second run was made with the initial detection made at a range predicted by convergence zone acoustic propagation. The probability of detection and range results from Section I were used for these two propagation paths.

The implemented target is radiating a 500 Hz tone at a source level, of 160 dB [1]. The transmission loss is computed using the sonar equation with spherical spreading out to a transition range, $R_0 = 3000$ yds, and cylindrical spreading beyond, plus a loss proportional to range. A noise background equivalent to that of the deep sea in sea state 3 on a quiet searching platform implies $NL = 66$ dB. The receiver characteristics are modeled using incoherent energy processing in a receiver band 100 Hz wide having a probability of detection equal to .5 with a .01 false alarm probability. The required receiver observation time under these constraints is 15 seconds. Using an absorption coefficient of .4 dB/kyd and the passive sonar equation, transmission loss = 86 dB.

A convergence zone with a 10 dB gain centered on 30 miles is assumed with an intermittent contact probability of .75. For ease of implementation a step approximation to the curve of Figure 4.1 is used with a worst case probability of .5 for the direct path case. The absorption coefficient for both the direct path and convergence zone cases is .4 dB/kyd. The expected direct path range of detection is 19 kyds. The resulting probability of detection as a function of range curve used in the simulations is depicted in Figure 4.1.




Figure 1

4.3 System Validation

The most encouraging result was that in all 54 simulation runs, the filter that was hypothesized to be the correct filter based on the chi-square statistic was either the "correct" filter, 000 degrees true at 6 knots, or an adjacent filter. The only qualification to this success was that the estimator was not always able to distinguish between the best direct path and convergence zone solutions. However, this inability to distinguish between the two propagation paths is a very real and important problem which is encountered in actual passive acoustic search cases.

The major problem with the estimator resulted from its instability at ranges less than about 3 n.m. At these relatively close ranges the non-linear aspects of the geometry tend to dominate and the linearized covariance prediction became inadequate. For this reason, most of the data for the direct path cases was taken after only 15-30 minutes into the run while ranges were still greater than 3 n.m. Although the convergence zone run data was all taken after 60 minutes into the run, the convergence zone had a much greater incidence of intermittent contact, receiving bearing information only about 75% of the time. Given these limitations, it is not surprising that the "correct" filters in each case were unable to significantly reduce their initial speed errors. In the direct path case, range errors were reduced in relatively short periods of time. However, in the convergence zone runs, with initial range errors as small as 3%, the estimator was only able to maintain the error margins approximately this small.

While clear differences in filter performance with respect to speed and range estimates were absent, there was a clear difference in filter estimates of target course for different initial target-searcher geometries. In particular, the searcher's estimates of target course during the beam aspect runs were consistently in error by 3-5 times the errors encountered with the other two initial geometries. This difference is again quite consistent with actual practice since a searcher presented the beam aspect of a target is in the worst possible location to detect relatively minor changes in course (or course errors).

II. Conclusion

The primary purpose of this project, to actually build a realistic ocean acoustic model and estimator for open-loop search strategy formulation, was accomplished. The various simulations which were run confirm the feasibility and performance of the system as a tool to be used in further strategy assessment. Even though any comprehensive strategy formulation would have to be based on using this system on a large number of Monte Carlo type runs, several strategy concepts suitable for a more global closed-loop implementation came to light with only a limited number of runs.

References

- [1] Urlick, R.J., "Principles of Underwater Sound", Second Edition, McGraw-Hill Book Company, New York, 1975.
- [2] Pedersen, M.A., "Comparison of Experimental and Theoretical Image Interference in Deep Water Acoustics", *J. Acoust. Soc. Am.*, 34:1197, 1962.
- [3] Knight, W.C., Fridham, R.G., and Kay, S.M., "Digital Signal Processing for Sonar", *Proceedings of the IEEE*, 69:2, November 1981, 1451-1506.
- [4] Weiss and Moor, "Improved Extended Kalman Filter Design for Passive Tracking", *IEEE Trans. on Auto. Control*, AC-25:4, August 1980, 807-811.
- [5] Petridis, V., "A Method for Bearings-Only Velocity and Position Estimation", *IEEE Trans Auto Contr.* AC-26:2, April 1981, 488-493.
- [6] Aidala, V.J., and Hammel, S.E., "Utilization of Modified Polar Coordinates for Bearings-Only Tracking", *IEEE Trans Auto Control*, AC-28:3, March 1983, 283-294.
- [7] Aidala, V.J., "Kalman Filter Behavior in Bearings-Only Tracking Applications", *IEEE Trans. Aerospace Electronic Systems*, AES-15, January 1979, 29-39.
- [8] Johnson, G.W., "Choice of Coordinates and Computational Difficulty", *IEEE Trans Auto Contr.* AC-19, February 1974, 77-80.
- [9] Van Trees, H.L., *Detection, Estimation, and Modulation Theory*, John Wiley & Sons, Inc., New York, N.Y. 1968, pp. 107-116.

DETECTION MODELING OF UNSATURATED AND PARTIALLY SATURATED OCEAN ACOUSTIC SIGNALS

Anastassios N. Perakis

Harilaos N. Psaraffis

Henry I. Gonzalez

Department of Naval Architecture
and Marine Engineering
The University of Michigan
Ann Arbor, Michigan 48109

Department of Ocean Engineering
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Summary

$N_X(n), N_Y(n)$ = zero-mean, uncorrelated Gaussian additive noise for the n th path

Furthermore, the envelope and the phase of the total signal are defined as:

$$\rho = (x^2 + y^2)^{1/2} \quad (2)$$

$$\phi = \tan^{-1}(y/x)$$

At short ranges and low frequencies, or for stable channels, the propagation is said to be unsaturated and the probability density function (PDF) of ρ is Rician and independent of the number of paths [1]. (In section 2 the distributions of ρ and its phase ϕ are presented).

At sufficiently long ranges and/or high frequencies, the propagation is fully saturated, which means that ϕ , the phase of ρ , is uniformly distributed between 0 and 2π , or each path has a phase θ_n that is normally distributed with a standard deviation $> 2\pi$. In this regime when $N > 4$ and the single path amplitudes r_n are approximately equal, phase random multipath propagation is obtained. It has been found [3] that the envelope ρ of a fully saturated phase random process obeys a Rayleigh PDF. Moreover, several other statistics and joint PDF's for the phase random process have been obtained, and are summarized in [6].

At intermediate ranges, where the signal experiences enough perturbations in the channel so that each θ_n can be characterized as a Gaussian random variable but with a standard deviation $< 2\pi$, partially saturated propagation is obtained. The frequency/range boundaries between the unsaturated, partially saturated, and fully saturated regimes are dependent upon the ocean dynamics or boundary dynamics of the propagation channel, as well as the magnitude of any relative source-receiver motion. Envelope statistics for signals in the partially saturated regime are presented in [1]. As the variance of the single path phase goes to zero, or becomes large, the first order PDF's converge to the unsaturated and fully saturated results respectively [1].

In previous publications [7,8] of the authors, continuous and discrete-time detection models using the results of phase random acoustic propagation [3,6] have been formulated. "Detection" was defined as an upcrossing of random variable ρ (the root mean square pressure at the passive sonar receiver) over a specified threshold ρ_0 . A continuous-time model was first developed for obtaining the PDF's of the time between two successive detections (interarrival time) and of the time between a detection and the first subsequent downcrossing through ρ_0 (holding time). The model was then compared with the extensively used (λ, σ) model and with available acoustic data. This model was seen to exhibit similar long-term behavior but markedly different short-term characteristics as compared with the (λ, σ) model, a fact which is due to the memory of the process. Comparison with data has demonstrated, in

The basic problem in ocean detection of narrowband acoustic signals is studied for unsaturated sound propagation, whereas considerable progress has been made towards modeling the more general case of partially saturated acoustic detection process. Detection is defined as occurring whenever ρ , the short-time average root mean square pressure at the receiver, exceeds a specified threshold level ρ_0 . In this paper, a two-state, discrete-time Markov model is first derived for the unsaturated ocean acoustic detection process. Closed-form expressions for the probability mass functions (PMF's) of the number of time steps separating either two successive detections (interarrival time) or one detection and the first subsequent "downcrossing" (holding time) are presented. Expressions for the joint probability density function (PDF) of ρ at two different points in time are obtained and used to determine the relevant one-step transition probabilities of the Markov model. Results using the model for various values of its input parameters are also presented and discussed. Another Markov model is next derived for partially saturated narrowband acoustic signal propagation, and closed-form expressions for the PMF's for the relevant interarrival and holding times are presented. The joint PDF of ρ at two different points in time is obtained from a rather general conditional PDF already derived by Middleton and checked to reduce to the limiting cases of the fully saturated and unsaturated processes. Future research may include a direct derivation of the partially saturated model, comparing both models with data available to us as well as with appropriate versions of the current memoryless "state-of-the-art" (λ, σ) model, and finally using each model in resource allocation schemes for target tracking, along the lines of previous research of the first two authors.

1. Introduction

In general the quadrature components of the envelope of a narrowband ocean acoustic multipath process are given by [1]

$$X = \sum_{n=1}^N (r_n \cos \theta_n + N_X(n)), \quad (1)$$

$$Y = \sum_{n=1}^N (r_n \sin \theta_n + N_Y(n)),$$

where

N = number of independent paths between source and receiver

r_n = the amplitude of the n th path

θ_n = the phase of the n th path

most cases, a significantly improved prediction capability over the (λ, σ) model.

Subsequently, a two-state and a four-state discrete-time Markov detection model were developed and closed-form expressions for the probability mass functions of the corresponding interarrival and holding times were derived. The results obtained using the latter models were favorably compared with both the continuous-time models and the data, the greatest improvement over the continuous-time models lying in the much lower computational effort involved.

In this paper, discrete-time detection models are developed for the unsaturated case, first for the "memoryless" case and then for the general "memory" Markov case. A discrete-time detection model is also proposed for the partially saturated acoustic propagation, which is shown to reduce to the two limiting cases of fully saturated and unsaturated processes. Conclusions and suggestions for future research form the final section of the paper.

2. Modeling the Unsaturated Detection Process

The probability density functions for the root-mean square pressure p and its phase ϕ for the unsaturated process are derived in Ref. 1. The density of p is Rician:

$$f_p(p) = \frac{p}{\sigma_N^2} \exp\left(-\frac{p^2 + R_S^2}{2\sigma_N^2}\right) \cdot I_0\left(\frac{pR_S}{\sigma_N^2}\right), \quad 0 \leq p < \infty \quad (3)$$

where

R_S = the magnitude of the constant signal vector

I_m = modified Bessel function of the first kind of order m

σ_N^2 = $N\sigma_{N_x}^2 = N\sigma_{N_y}^2$, where

$\sigma_{N_x}^2, \sigma_{N_y}^2$ = the variances of $N_x(n)$ and $N_y(n)$, respectively, as defined in the introduction

In our previous Markov modeling of the phase random process (fully saturated sound propagation), a two-state model and a four state model were developed. [7,8] Comparison with data [8] has revealed that both models, when properly calibrated, yield very satisfactory results, the two-state being consistently as accurate as the four-state model.

We will henceforth restrict ourselves in developing a two-state Markov model for the unsaturated process of the general form shown in Figure 1.




Figure 1. Two-state discrete-time Markov model

where "U" = "up" state, defined by $p > p_0$
 "D" = "down" state, defined by $p < p_0$
 a = prob $(p_2 < p_0 | p_1 > p_0)$

(4)

$$b = \text{prob}(p_1 > p_0 | p_2 < p_0) \quad (5)$$

In the memoryless case, $a + b = 1$, and

$$a = \text{prob}(p_2 < p_0 | p_1 > p_0) = \text{prob}(p_2 < p_0 | p_1 < p_0) = \text{prob}(p_2 < p_0)$$

$$\text{or } a = \int_0^{p_0} f_p(p) dp = 1 - Q\left(\frac{R_S}{\sigma_N}, \frac{p_0}{\sigma_N}\right) \quad (6)$$

with $f_p(p)$ as in (3),

and the generalized O-function defined in [5] as

$$O_M(a, b) = \int_0^\infty \xi \left(\frac{\xi}{a}\right)^{M-1} \exp\left(-\frac{a^2 + \xi^2}{2}\right) I_{M-1}(a\xi) d\xi \quad (7)$$

$M = 1, 2, \dots$

From (6), we can proceed to evaluate the probability mass functions for the interarrival and holding time. In general, these PMF's take the form [8]

$$P_H(k) = (1-a)k^{-1}a, \quad k = 1, 2, \dots (\text{holding time}) \quad (8)$$

$$P_I(n) = \frac{ab}{a-b} [(1-b)^{n-1} - (1-a)^{n-1}], \quad n = 2, 3, \dots (\text{interarrival time}) \quad (9)$$

In the memoryless case, (9) becomes

$$P_I(n) = \frac{a(1-a)}{(1-2a)} [(1-a)^{n-1} - a^{n-1}], \quad n = 2, 3, \dots \quad (10)$$

In the non-trivial case ($a+b \neq 1$) the calculation of the transition probabilities requires knowledge of the joint density function $f_{p_1 p_2}(p_1, p_2)$. This second-order density has already been derived in a rather general form by Middleton's [9] treatment of the statistical properties of additive narrowband signal and normal noise processes.

Using Middleton's results and after extensive algebraic manipulations, we obtain:

$$f_{p_1 p_2}(p_1, p_2) = \frac{p_1 p_2}{\sigma_N^4 (1-r_0^2)} \exp\left(-\frac{p_1^2 + p_2^2}{2\sigma_N^2 (1-r_0^2)}\right) \exp\left(\frac{A_0^2}{\sigma_N^2 (1+r_0)}\right) \cdot \sum_{m=0}^{\infty} \epsilon_m I_m\left(\frac{r_0 p_1 p_2}{\sigma_N^2 (1-r_0^2)}\right) I_m\left(-\frac{A_0 p_1}{\sigma_N^2 (1+r_0)}\right) I_m\left(\frac{A_0 p_2}{\sigma_N^2 (1+r_0)}\right) \quad (11)$$

where $\epsilon = 1, -1, 2, -2, \dots$

It is reasonable to expect that (11) will reduce, for $t \rightarrow \infty$, to the product of $f_{p_1}(p_1) \cdot f_{p_2}(p_2)$. Since $r_0 \rightarrow 0$ for $t \rightarrow \infty$ (uncorrelatedness), (11) gives:

$$f_{p_1 p_2}(p_1, p_2) = \frac{p_1 p_2}{\sigma_N^4} \exp \left(- \frac{p_1^2 + p_2^2 + 2\lambda_0^2}{2\sigma_N^2} \right)$$

$$I_0 \left(\frac{\lambda_0 p_1}{\sigma_N^2} \right) I_0 \left(\frac{\lambda_0 p_2}{\sigma_N^2} \right) \quad (12)$$

$$\text{nce } I_m(x) = \frac{(x/2)^m}{m!} \quad \begin{matrix} 0, m \neq 0 \\ 1, m = 0 \end{matrix} \quad (13)$$

(12) can be rewritten as

$$f_{p_1 p_2}(p_1, p_2) = \frac{p_1}{\sigma_N^2} \exp \left(- \frac{p_1^2 + \lambda_0^2}{2\sigma_N^2} \right) I_0 \left(\frac{\lambda_0 p_1}{\sigma_N^2} \right) \frac{p_2}{\sigma_N^2} \exp \left(- \frac{p_2^2 + \lambda_0^2}{2\sigma_N^2} \right) I_0 \left(\frac{\lambda_0 p_2}{\sigma_N^2} \right) \quad (14)$$

$$\text{e. } f_{p_1 p_2}(p_1, p_2) + f_{p_1}(p_1) + f_{p_2}(p_2) \quad (15)$$

t =

we can now proceed to evaluate the one-step transition probabilities of the Markov model.

$$12 = a = \int_0^{p_0} \int_0^{p_0} f_{p_1 p_2}(p_1, p_2) dp_1 dp_2 / \int_0^{p_0} f_p(p) dp \quad (16)$$

$$21 = b = \int_0^{p_0} \int_0^{p_0} f_{p_1 p_2}(p_1, p_2) dp_1 dp_2 / \int_0^{p_0} f_p(p) dp \quad (17)$$

$$\text{and } P_{11} = 1 - P_{12}, P_{22} = 1 - P_{21} \quad (18)$$

The double integrals in (16) and (17) can be evaluated as functions of

$$\lambda_1 = \int_0^{p_0} \int_0^{p_0} f_{p_1 p_2}(p_1, p_2) dp_1 dp_2 \quad (19)$$

Although λ_1 is symmetric with respect to p_1 and p_2 , it cannot be expressed as a product of one function of p_1 and one of p_2 . Instead, we can rewrite (20), taking (11) into account, as follows:

$$\lambda_1 = \int_0^{p_0} \frac{p_1}{\sigma_N^4(1-p_0^2)} \exp \left(- \frac{p_1^2}{2\sigma_N^2(1-p_0^2)} - \frac{\lambda_0^2}{\sigma_N^2(1+p_0)} \right) \cdot \int_0^{p_0} p_2 \exp \left(- \frac{p_2^2}{2\sigma_N^2(1-p_0^2)} \right) \sum_{m=0}^{\infty} I_m \left(\frac{\lambda_0 p_1 p_2}{\sigma_N^2(1-p_0^2)} \right) I_m \left(\frac{\lambda_0 p_1}{\sigma_N^2(1+p_0^2)} \right) I_m \left(\frac{\lambda_0 p_2}{\sigma_N^2(1+p_0^2)} \right) dp_2 dp_1 \quad (20)$$

In the above, λ_0 is identical to R_S of eqn. (3).

Having evaluated λ_1 , eqns. (16) and (17) can be expressed as:

$$P_{12} = \frac{P_u - \lambda_1}{1 - P_u} \quad (21)$$

$$P_{21} = 1 - \frac{\lambda_1}{P_u} \quad (22)$$

$$P_{11} = 1 - P_{12}, P_{22} = 1 - P_{21} \quad (23)$$

where P_u is the (unconditional) probability of p being less than p_0 . Efforts to simplify the evaluation of λ_1 in Eq. (20) were not successful. The double numerical integration of a function involving the infinite sum of products of three modified Bessel functions was expected to and did actually produce computational problems (excessive CPU time). These were partially alleviated using the asymptotic properties of the Bessel functions involved, in determining the tolerances employed in terminating the evaluation of the infinite summations. Still, for extreme (that is, too small or too large) detection thresholds, the computational effort is unacceptably large. However, this is not expected to be a problem in practice, since we do not need to use such extreme thresholds - in fact, they result in memoryless Markov Models and the problem does not exist, since the evaluation of the relevant one-step transition probabilities requires the knowledge of just the unconditional distributions.

Implementing the Model

The Markov model for the unsaturated detection process was used with a variety of - hopefully appropriate - inputs for the parameters involved, namely σ_N^2 , v , R_S , T and p_0 . No comparisons of our predictions with real acoustic data that could be appropriately modeled as unsaturated sound propagation are presented. However, we do have such data at our disposal and we are planning to use them to that effect (see section 4).

Figs. 3-4 present typical results using $\sigma_N^2 = 1.58$, $R_S = 2.23$, ($R_S^2 = 5$), $v = .2\text{Hz}$, a time step of .4 sec and thresholds $p_0 = 1.58, 2.37$, and 3.16 . In Figs. 3 and 4, and for the $p_0 = 3.16$ threshold only, both the histogram and the "outline" (i.e., the continuous line passing through the top midpoint of each bar in the histogram) of the corresponding PMP were drawn. For the other two thresholds, only the "outlines" were plotted, since including their histograms would most probably reduce the legibility of these figures. Of these values, σ_N^2 and R_S were picked from an unsaturated example [1] and the rest were chosen by the authors and are more or less arbitrarily. It is clear that, although the density of the holding time is very sensitive to the magnitude of the detection threshold, the density of the interarrival time is much less threshold dependent. This reminds us of what we would get in a pure sinusoidal signal situation, where we have a constant interarrival time (equal to the period of the sinusoidal signal) but different holding times for each threshold (Figure 2).

Such a result was not observed in our previous study of the detection process using the phase random model for ocean propagation [8]. In the unsaturated case, $p(t)$ is obviously not a strict sinusoid and hence we do not get the above δ -function densities for the interarrival and holding times. Still, Fig. 4 shows that the timing of detection events (i.e., the distribution of the interarrival times) is almost independent of the detection threshold. This threshold makes its presence felt only in the distributions of the holding time, in which we quite obviously have shorter holding times for higher thresholds.

Figures 5 and 6 demonstrate the relative insensitivity of the above results to changes in the time step T . Figs. 9 and 10 correspond to $T = .8$ sec.,

or twice that of Figs. 3 and 4 with which they have all other inputs in common.




Figure 2. Sinusoidal Signal, Interarrival and Holding Time PDF's




Figure 3. Holding Time, $\rho_0 = 1.58, 2.23, 3.16$, $T = .4$ sec., $R_S = 2.23$




Figure 4. Interarrival Time, $\rho_0 = 1.58, 2.23, 3.16$, $T = .4$ sec., $R_S = 2.23$

Finally, Figs. 7 and 8 present the results obtained using the unsaturated model with $R_S = 0.0$ and other inputs as in Figs. 3-4. Typical unsaturated propagation involves a large R_S (constant vector) and additional Gaussian noise, and the $R_S = 0$ case could be perhaps more "naturally" described by a phase random model.




Figure 5. Holding Time, $\rho_0 = 3.16$, $T = .4$ sec., $R_S = 2.23$




Figure 6. Interarrival Time, $\rho_0 = 3.16$, $T = .4$ sec., $R_S = 2.23$




Figure 7. Holding Time, $\rho_0 = 1.58$, $T = .8$ sec., $R_S = 2.23$




Figure 8. Interarrival Time, $\rho_0 = 1.58$, $T = .8$ sec., $R_S = 2.23$

3. Modeling the Partially Saturated Detection Process

Following the notation of the previous section, the first-order PDF of ρ for a partially saturated process has been derived [1] and is of the form:

$$f_{\rho}(\rho) = \frac{\rho}{2\pi\sigma_x\sigma_y(1-\rho_{xy}^2)^{1/2}} \left| \exp \frac{-1}{2(1-\rho_{xy}^2)} \right. \\ \left. \frac{\rho^2 + 2\rho_x^2}{2x} - \frac{2\rho_{xy}\rho_x\rho_y}{\sigma_x\sigma_y} + \frac{\rho^2 + 2\rho_y^2}{2y^2} \right| \times \\ \int_0^{2\pi} \exp \frac{-1}{2(1-\rho_{xy}^2)} [a \cos(2\phi) + b \sin(2\phi) + \\ c \cos(\phi) + d \sin(\phi)] d\phi \quad (24)$$

where:

$$a = \rho^2(1/2\sigma_x^2 - 1/2\sigma_y^2)$$

$$b = -\rho^2 \frac{\rho_{xy}}{\sigma_x\sigma_y}$$

$$c = \rho \frac{2\rho_{xy}\rho_y}{\sigma_x\sigma_y} - \frac{2\rho_x}{\sigma_x^2}$$

$$d = \rho \frac{2\rho_{xy}\rho_x}{\sigma_x\sigma_y} - \frac{2\rho_y}{\sigma_y^2}$$

For a memoryless Markov model for the unsaturated process, (24) is adequate. for the general (memory) case, however, we need to know the joint PDF $f_{\rho_1\rho_2}(\rho_1, \rho_2)$. This is given by

$$f_{\rho_1\rho_2}(\rho_1, \rho_2) = \int_{\phi} f_{\rho_1\rho_2}(\rho_1, \rho_2 | \phi) f_{\phi}(\phi) d\phi, \text{ where}$$

$f_{\rho_1\rho_2}(\phi)$ conditioned upon a general set of parameters is obtained in [9] as:

$$f_{\rho_1\rho_2}(\rho_1, \rho_2 | \phi) = \frac{\rho_1\rho_2}{\sigma_N^4(1-k_0^2)} \exp \frac{-(\rho_1^2 + \rho_2^2)}{2\sigma_N^2(1-k_0^2)}$$

$$B_2(A_1, A_2, \epsilon_1, \epsilon_2) \propto \epsilon_m \int \frac{k_0\rho_1\rho_2}{\sigma_N^2(1-k_0^2)}$$

$$I_1 \frac{Y_{12}\rho_1}{\sigma_N^2(1-k_0^2)} T_m \frac{Y_{21}\rho_2}{\sigma_N^2(1-k_0^2)} \cos m(\psi_{21} - \psi_{12} + \phi_0) \quad (25)$$

with ϕ not explicitly appearing on the right hand side, and

$$-Y_{12} = [A_1^2 + k_0^2 A_2^2 - 2k_0 A_1 A_2 \cos(\epsilon_2 - \epsilon_1 + \phi_0)]^{1/2}$$

$$-\psi_{12} = \tan^{-1} \left| \frac{k_0 A_2 \sin(\epsilon_2 + \phi_0) - A_1 \sin \epsilon_1}{A_1 \cos \epsilon_1 - k_0 A_2 \cos(\epsilon_2 + \phi_0)} \right|$$

$-\psi_{21}$ and Y_{21} are obtained by interchanging subscripts 1 and 2

$$-B_2(A_1, A_2, \epsilon_1, \epsilon_2) = \exp[-(A_1^2 + A_2^2 - 2k_0 A_1 A_2 \cos(\epsilon_2 - \epsilon_1 + \phi_0))/2\sigma_N^2(1 - k_0^2)]$$

$$-\phi_0 = \tan^{-1}[\lambda_0(t)/\rho_0(t)]$$

$$-k_0 = (r_0^2(t) + \lambda_0^2(t))^{1/2}$$

$$-r_0(t) = E(N_{C_1} N_{C_2})/\sigma_N^2$$

$$-\lambda_0(t) = E(N_{C_1} N_{S_2})/\sigma_N^2$$

$$-\epsilon_n = \omega_d t - \psi(t; \phi)$$

$$-\omega_d = \omega_c - \omega_0$$

Since a partially saturated process essentially consists of an additive narrowband signal with Gaussian noise superimposed, (25) should serve as our second-order density [12].

Necessary conditions for this to be true is that the second-order partially saturated density should reduce to the densities of the limiting cases of fully saturated and unsaturated processes.

Three assumptions are necessary to reduce Eq. (2) to the unsaturated case: 1. the signal is an ensemble of unmodulated sinusoids (then $A_1 = A_2 = \lambda_0$), 2. $f_c = f_0$, (thus $\omega_d = 0$), 3. the noise spectrum is symmetrical about $f_c = f_0$. It then follows that:

$$\lambda_0(t) = 0, \quad k_0 = r_0(t), \text{ and } \phi_0 = 0$$

and the following simplifications occur:

$$\epsilon_1 = \epsilon_2$$

$$B_2 = \exp(-\lambda_0^2/[\sigma_N^2(r_0 + 1)])$$

$$\psi_{12} = \psi_{21}$$

$$Y_{12} = Y_{21} = [\lambda_0^2(1 + r_0^2 - 2r_0)]^{1/2} = \lambda_0(1 - r_0)$$

Finally, (26) becomes:

$$f_{\rho_1\rho_2}(\rho_1, \rho_2) = [\rho_1\rho_2/\sigma_N^4(1-r_0^2)] \\ \exp(-\rho_1^2 + \rho_2^2/[2\sigma_N^2(1-r_0^2)]) \exp(-\lambda_0^2/[\sigma_N^2(r_0+1)]) \times \\ \int \epsilon_m \int I_m(r_0\rho_1\rho_2/\sigma_N^2(1-r_0^2)) I_m(\lambda_0\rho_1/\sigma_N^2(r_0+1)) \\ I_m(\lambda_0\rho_2/\sigma_N^2(r_0+1)) \quad (26)$$

(26) is the second-order PDF for an unsaturated process, coinciding with (26) of [12].

For reduction to the fully saturated case, two assumptions are necessary. Since observed S/N ratios for fully saturated samples are very low, the signal

can be modelled as pure noise. Thus $A_1 = A_2 = 0$. Also, the noise spectrum is symmetrical about $f_c = f_0$, hence

$$\lambda_0(t) = 0, \quad k_0 = r_0(t), \quad \text{and} \quad \phi_0 = 0.$$

It then follows that:

$$\epsilon_1 = \epsilon_2, \quad u_d = u_c, \quad B_2 = 1, \quad \text{and} \quad Y_{12} = Y_{21} = 0.$$

And since $I_m(0) = 0 \neq 0$, $I_m(0) = 1$, $m = 0$, the second-order PDF, Equation (2) becomes:

$$f_{\rho, \rho_2}(\rho_1, \rho_2) = [\rho_1 \rho_2 / \sigma_N^4 (1 - r_0^2)] \exp(-[\rho_1^2 + \rho_2^2] / [2\sigma_N^2 (1 - r_0^2)]) I_0[\rho_1 \rho_2 / \sigma_N^2 (1 - r_0^2)] \quad (27)$$

Equation (27) is the second-order PDF for a fully saturated process, as given in equations (9.20) of [9] and (62) of [3].

Thus, we have met necessary conditions for (25) to be the second-order PDF for a partially saturated process. It remains to be seen whether this is also a sufficient condition.

Another necessary condition is that the second-order PDF, Equation (2), must reduce to the first-order PDF, (24), when integrated over ρ_1 or ρ_2 .

4. Conclusions and Future Research

In this paper, an analytical model for the unsaturated acoustic detection process was presented and probability mass functions for the interarrival and holding times were derived. The unsaturated mode of acoustic propagation was seen to exhibit different characteristics than the previously developed phase random acoustic detection models. A major difference between these two modes of acoustic propagation lies in the narrower (for the unsaturated case) distribution of ρ , which approaches a normal density as R_g grows large. A more striking difference lies in the relative independence of the interarrival time MF to the detection threshold ρ_0 for the unsaturated case, a property not observed in the fully saturated models.

In the third section of this paper, a discrete-time Markov model for the partially saturated detection process was presented. The second-order PDF of ρ has been conjectured to be one presented in [9], (appropriately integrated over the conditioning set of parameters ϕ). It was also shown that this PDF reduces, as required, to the limiting cases of fully saturated and unsaturated propagation. Clearly, once the joint PDF is obtained, the PDFs of the interarrival and holding times can be readily derived.

Several interesting extensions of the present research could be pursued, the most urgent of which might be the comparison of both unsaturated and partially saturated models with those acoustic data records in our disposal for which the PDF of ρ approaches the Rician or the partially saturated density respectively.

A relevant problem in this comparison is to devise a method for estimating the parameters of these distributions from the data (in the phase-random case, v and σ_1^2 were estimated by fitting a normal distribution to the histogram of Λ , where $\Lambda \approx 10 \log_{10} \rho^2$ and $\Lambda = dA/dt$). Once and if the ability of our models to predict the timing of detection events is demonstrated, we can very conveniently employ them in all sequential optimization algorithms appearing in [2, 10, 11] for a more rational allocation of acoustic sensor resources in target tracking, and also in the

sequential hypothesis testing algorithms of [10] that will help us in signal vs. noise, target vs. several false targets, data association, target identification and other relevant problems. No changes need to be made in the algorithms of [10] other than substituting the phase random model for acoustic propagation with the unsaturated model developed in this paper.

Acknowledgments

The work has been partially supported by the Office of Naval Research, Contract No. N00014-79-C-0239. A summary of the unsaturated detection modelling will appear as a correspondence in the Journal of the Acoustical Society of America. We are indebted to Lieut. Dr. Peter Mikhalevsky for providing us with the first order statistics and for his comments.

References

- [1] Mikhalevsky, P.N., "Envelope Statistics of Partially Saturated Processes," J. Acoust. Soc. Am. 72(1), pp. 151-158, July 1982.
- [2] Psaraftis, N.N. and Perakis, A.N., "Myopic and Preshvopic Approaches to a Multi-Sensor, Multi-Target Resource Allocation Problem," 5th MIT/ONR Symposium on Command, Control and Communications, Monterey, CA, August 1982.
- [3] Hamblen, W.R., "A Phase Random Multipath Model for Acoustic Signal Fluctuation in the Ocean and its Comparison with Data," Ph.D. thesis, Massachusetts Institute of Technology, 1978.
- [4] Perakis, A.N. and N.N. Psaraftis, "Discrete-Time Detection Modeling for Unsaturated Ocean Propagation," Working Paper No. OE-ONR-82-2, October 1982.
- [5] Rappaport, S.A., "Computing Approximations to the Generalized O-Function and its Complement," Correspondence, T.F.E.E. Transactions on Information Theory, July 1971, pp. 497-498.
- [6] Mikhalevsky, P.N., "First Order Statistics for finite bandwidth multi-path signals with and without frequency or phase modulation," J. Acoust. Soc. Am. 66(3), pp. 751-762, 1979.
- [7] Psaraftis, N.N., Perakis, A.N. and Mikhalevsky, P.N., "New Models for the Ocean Acoustic Detection Process," J. Acoust. Soc. Am., 69(6), 1724-1734, June 1981.
- [8] Psaraftis, N.N., Perakis, A.N., and Mikhalevsky, P.N., "Memory Detection Models for Phase-Random Ocean Acoustic Fluctuations," International Conference on Communications, Denver, CO, June 1981.
- [9] Middleton, David, "An Introduction to Statistical Communication Theory," McGraw Hill, 1960.
- [10] Perakis, A.N., "Contributions to Resource Allocation in Target Tracking Using New Ocean Acoustic Detection Models," Ph.D. Thesis, O.E., M.I.T. 1982.
- [11] Psaraftis, N.N. and A.N. Perakis, "A Basic Problem of Resource Allocation in Target Tracking," J. Acous. Soc. Am. 72(3), pp. 824-833, September 1982.
- [12] Gonzalez, H.I. and N.N. Psaraftis, "Detection Models for Partially Saturated Ocean Acoustic Propagation," working paper DE-ONR-83-1, MIT, Feb. 1983.

SECTION IV

Decision: Organizations and Theory

DISTRIBUTED OPTIMIZATION ALGORITHMS WITH COMMUNICATIONS*

by

John Tsitsiklis, Dimitri P. Bertsekas and Michael Athans
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

ABSTRACT

We discuss the convergence properties of asynchronous distributed iterative optimization algorithms, tolerating communication delays. We focus on a gradient-type algorithm for minimizing an additive cost function and present sufficient conditions for convergence. We view such an algorithm as a model of adjustment of the decisions of decision makers in an organization and we suggest that our results can be interpreted as guidelines for designing the information flows in an organization.

1. Introduction

This paper concerns the convergence properties and communication requirements of asynchronous distributed optimization algorithms, tolerating communication delays. The results being presented may be interpreted as pertaining to the performance of potential parallel computing machines. Alternatively, an approach which we pursue in this paper, our results may be viewed as a description of the adjustment process in a distributed organization, possibly involving human decision makers. Moreover, it could be maintained that the mathematical models discussed here, capture some aspects of the ever-present "bounded rationality" of human decision makers [Simon, 1980].

Our motivation is the following: A boundedly rational decision maker solving an optimization problem (minimize $J(x)$), may be viewed as an iterative optimization algorithm, whereby a tentative decision $x(n)$ is made at time n , and then the decision is updated, in a direction of improvement. For example, we may have

$$x(n+1) = x(n) - \gamma \frac{\partial J}{\partial x}(x(n)), \quad (1.1)$$

which corresponds to the well-known gradient algorithm. By extending the above analogy to more complex settings, an organization (or, at least, some aspects of it) consisting of cooperative, boundedly rational decision makers may be viewed as a distributed algorithm. For example, suppose that x is a decision vector and that the i -th decision maker is in charge of the i -th component x_i of x , which he updates according to

$$x_i(n+1) = x_i(n) - \gamma_i \frac{\partial J}{\partial x_i}(x(n)). \quad (1.2)$$

If each decision maker was to update his part of the decision (his own component), at each instance of time according to (1.2), we would have a synchronous distributed implementation of the centralized gradient algorithm. Synchronous algorithms have been studied in a variety of contexts [Arrow and Hurwicz, 1960; Gallager, 1977] but they also have certain drawbacks: (1) Decision maker i , in order to update $x_i(n)$ according to (1.2), he needs to know $x(n)$, at time n . This requires that each decision maker informs all others,

at each time instance, on the adjustments of his decisions. So, in some sense, the synchronous model requires "a lot of communications."

(2) A second drawback of synchronous algorithms is that communication delays can introduce bottlenecks and slow down the algorithm. In particular, the time between two consecutive updates has to be at least as large as the maximum communication delay between any pair of decision makers.

(3) Finally, complete synchronization is certainly an unrealistic model of human organizations.

For the above reasons, we choose to study asynchronous distributed versions of iterative optimization algorithms, in which decision makers do not need to communicate to every other decision maker at each time instance. Such algorithms avoid communication overloads, they are not excessively slowed down by communication delays and there is not even a requirement that each decision maker updates his decision at each time instance, which makes them even more realistic.

2. General Properties and Convergence Conditions of Asynchronous Distributed Algorithms

We now discuss the main principle underlying the class of asynchronous algorithms which we consider: as we mentioned, in Section 1, for a synchronous algorithm, each decision maker needs to be informed of the most recent value of the decisions of all other decision makers. Suppose now that decision maker i , at time n , needs for his computations the current value $x_j(n)$

of the j -th component of x , but he does not know this value. We then postulate that decision maker i will carry out his computations as in the synchronous algorithm, except that (not knowing $x_j(n)$) he will use the value of x_j in the most recent message he has received from decision maker j . Due to asynchronism and communication delays, decision maker i will, in general, use out-dated values of x_j to update his own decisions.

However, updates based on out-dated information may be substantially better than not updating at all. The crucial questions which arise are: How much out-dated information may be tolerated? How frequent should communications be, so that the distributed algorithm operates in a desirable manner?

Questions of this nature have been addressed by Bertsekas [1982, 1983] for the distributed version of the successive approximations algorithm for dynamic programming and the distributed computation of fixed points. We have obtained general convergence results of a related nature for the asynchronous distributed versions of deterministic and stochastic iterative pseudo-gradient [Poljak and Tsytkin, 1973] (or "descent-type") algorithms. Some representative algorithms covered by our general results are deterministic

* Research supported by ONR under Contract ONR/N00014-77-C-0532 (NR-041-519).

Equations (3.8), (3.9) together with

$$x_i^i(n+1) = x_i^i(n) - \gamma_i \sum_{j=1}^M \lambda_{ij}^j(n) \quad (3.10)$$

specify completely the asynchronous distributed algorithm of interest.

Let us now assume that the time between consecutive communications and the communication delays are bounded. We allow, however, these bounds to be different for each pair of processors and each type of message:

Assumption: For some constants P^{ik}, Q^{ik} ,

$$n - P^{ik} \leq p^{ik}(n) \leq n, \quad \forall (i,k) \in E, \forall n, \quad (3.11)$$

$$n - Q^{ik} \leq q^{ik}(n) \leq n, \quad \forall (k,i) \in E, \forall n, \quad (3.12)$$

Note that we may let $P^{ii} = Q^{ii} = 0$.

The following result states that the algorithm converges if P^{ik} and Q^{ik} are not too large compared to the degree of coupling between different subproblems. [Tsitsiklis, 1983].

Theorem 3.1: Suppose that for each i

$$\frac{2}{\gamma_i} > \sum_{j=1}^M K_{ij}^{ij} + \sum_{k=1}^M \sum_{j=1}^M K_{ij}^{kj} (P^{ik} + Q^{kj} + P^{jk} + Q^{ki}). \quad (3.13)$$

Let $z(n) = (x_1^1(n), x_2^2(n), \dots, x_M^M(n))$. Then,

$$\lim_{n \rightarrow \infty} \frac{\partial J}{\partial x_i} (z(n)) = 0, \quad \forall i, j. \quad (3.14)$$

We close this section with a few remarks:

1. The bounds provided by (3.13) are sufficient for convergence but not necessary. It is known [Nertsekas, 1983] that a decentralized algorithm of a similar type may converge in certain special cases, even if the i 's are held fixed, while the bounds P^{ik}, Q^{ik} are allowed to be arbitrarily large. So, the gap between the sufficient conditions (3.13) and the necessary conditions may be substantial. Further research should narrow this gap.
2. The convergence rate of the distributed algorithm should be expected to deteriorate as the bounds P^{ik}, Q^{ik} increase. A characterization of the convergence rate, however, seems to be a fairly hard problem.

4. Towards Organizational Design

Suppose that we have a divisionalized organization and that the objective of the organization is to minimize a cost J which is the sum of the costs J^i faced by each division. To each division, there corresponds a decision maker which is knowledgeable enough about the structure of the problem he is facing, to the extent that given a tentative decision he is able to change his decision in a direction of improvement. Moreover, suppose that the divisions are interacting in some way; that is, the decision of one decision maker may affect the costs of another division. Suppose, finally, that decision makers regularly update their decisions taking into account the decisions of other decision makers and the effects of their own decisions on other divisions. Messages are being exchanged from time to time carrying the required information. Clearly, the mathematical model of Section 3 may be viewed as a model of the above situation.

A natural question raised by the above described situation concerns the design of the information flows within the organization, so as to guarantee smooth

operation. But this is precisely the issue addressed by Theorem 3.1: the bounds K_{ij}^k may be thought as quantifying the degree of coupling between divisions; the bounds P^{ij}, Q^{ij} describe the frequency of communications and γ_i represents the speed of adjustment. Theorem

3.1 links all these quantities together and provides some conditions for smooth operation, whereby communication rates are prescribed in terms of the degree of coupling.

We may conclude that the approach of Section 3 may form the basis of a procedure for designing an organizational structure, or -more precisely- the information flows within an organization. Of course, Theorem 3.1 does not exhaust the subject. In particular, Theorem 3.1 suggests a set of feasible organizational structures, with generally different convergence rates. There remains the problem of choosing a "best" such structure.

It is also conceivable that the structure of the underlying optimization problem slowly changes with time, and so do the bounds K_{ij}^k , but in a time scale slower than the time scale of the adjustment process.

In such a case, the bounds P^{ij}, Q^{ij} should also change. This leads to a natural two-level organizational structure: At the lower level, we have a set of decision makers continuously adjusting their decisions and exchanging messages. At a higher level, we have a supervisor who monitors changes in K_{ij}^k and accordingly instructs the low-level decision makers to adjust their communication rates. Note that the supervisor does not need to know the details of the cost function; he only needs to know the degree of coupling between divisions. This seems to reflect the actual structure of existing organizations. Low level decision makers are "experts" on the problems facing them, while higher level decision makers only know certain structural properties of the overall problem and make certain global decision, e.g. setting the communication rates.

Event-Driven Communications

We now discuss a slightly different "mode of operation" for the asynchronous algorithm, which has also clear organizational implications. It should be clear that communications are required by the distributed algorithm so that decision makers are informed of changes occurring elsewhere in the system. Moreover, the bounds P^{ik}, Q^{ik} of Section 3 effectively guarantee that a message is being sent whenever a substantial change occurs. The same effect, however, could be accomplished without imposing bounds on the time between consecutive message transmissions: each decision maker could just monitor his decisions and inform the others whenever a substantial change occurs. It seems that the latter approach could result in significant savings in the number of messages being exchanged, but further research is needed on this topic.

5. Conclusions

A large class of deterministic and stochastic iterative optimization algorithms admit natural distributed asynchronous implementations. Such implementations (when compared to their synchronous counterparts) may retain the desired convergence properties, while reducing communication requirements and removing bottlenecks caused by communication delays.

We have focused on a deterministic gradient-type algorithm for an additive cost function and we have shown that the communication requirements depend in a natural way on the degree of coupling between different components of the cost function. This approach addresses the basic problem of designing the information flows

gradient-type algorithms, as well as stochastic approximations algorithms. Due to space considerations, we only discuss here the nature of the results. Exact statements and the proofs may be found in [Tsitsiklis, 1983] and in forthcoming publications. Preliminary versions of these results appear in [Tsitsiklis, Bertsekas and Athans, 1983].

To discuss the nature of the convergence conditions, we distinguish two cases:

A. Constant Step-Size Algorithms (e.g. gradient algorithm)

For such algorithms it has been shown that convergence to the centralized optimum is obtained, provided that the time between consecutive communications between pairs of decision makers, plus the communication delay, is bounded by an appropriate constant. Moreover, the larger the step-size (i.e. the constant γ in equation (1.2)), the smaller the above mentioned constant. The latter statement admits the appealing interpretation that the larger the updates by each decision maker, the more frequent communications are required.

B. Decreasing Step-Size Algorithms (e.g. stochastic approximation algorithms)

In this case, the algorithm becomes slower and slower as the time index increases. This allows the process of communications to become progressively slower, as well. In particular, it has been shown that convergence to the centralized optimum is obtained even if the time between consecutive communications between pairs of decision makers, plus communication delays, increase without bound, as the algorithm proceeds, provided that the rate of increase is not too fast.

3. A Distributed Gradient Algorithm

In this section we consider a rather simple distributed algorithm for minimizing an additive cost function. Due to the simplicity of the algorithm, we are able to derive convergence conditions which are generally tighter than the general conditions discussed in the previous sections. It will be seen shortly, that these conditions admit appealing organizational interpretations.

The conceptual motivation behind our approach is based on the following statement:

If an optimization problem consists of subproblems, each subproblem being assigned to a different decision maker, then the frequency of communications between a pair of decision makers should reflect the degree by which their subproblems are coupled.

The above statement is fairly hard to capture mathematically. This is accomplished, however, to some extent, by the model and the results of this section.

Let $J: R^M \rightarrow R$ be a cost function to be minimized with a special structure:

$$J(x) = J(x_1, \dots, x_M) = \sum_{i=1}^M J^i(x_1, \dots, x_M) \quad (3.1)$$

where $J^i: R^M \rightarrow R$. So far, equation (3.1) does not impose any restriction on J ; we will be interested, however, in the case where, for each i , J^i depends on x_i and only a few more components of x ; consequently, the Hessian matrix of each J^i is sparse.

We view J^i as a cost directly faced by the i -th decision maker. This decision maker is free to fix or update the component x_i , but his cost also depends on a few interaction variables (other components of x)

which are under the authority of other decision makers.

We may visualize the structure of the interactions by means of a directed graph $G=(V,E)$:

- (i) The set V of nodes of G is $V=\{1, \dots, M\}$
- (ii) The set of edges E of the graph is

$$E=\{(i,j): J^j \text{ depends on } x_i\} \quad (3.2)$$

Since we are interested in the fine structure of the optimization problem, we quantify the interactions between subproblems by assuming that the following bounds are available:

$$\left| \frac{\partial^2 J}{\partial x_i \partial x_j} \right| \leq K_{ij}, \quad \left| \frac{\partial^2 J^k}{\partial x_i \partial x_j} \right| \leq K_{ij}^k, \quad \forall x \in R^M, \quad (3.3)$$

where (without loss of generality)

$$K_{ij} \leq \sum_{k=1}^M K_{ij}^k. \quad (3.4)$$

A synchronous distributed gradient-type algorithm for this problem could be:

1. For each $(i,j) \in E$, decision maker j evaluates

$$\lambda_{ij}^j(n) = \frac{\partial J^j}{\partial x_i}(x(n)) \quad (3.5)$$

2. For each $(i,j) \in E$, decision maker j transmits $\lambda_{ij}^j(n)$ to decision maker i .
3. Each decision maker i updates x_i according to

$$x_i(n+1) = x_i(n) - \gamma_i \sum_{j=1}^M \lambda_{ij}^j(n) \quad (3.6)$$

4. For each $(i,j) \in E$, decision maker i transmits $x_i(n+1)$ to decision maker j .

We now consider the asynchronous version of the above algorithm. Let $x^i(n) = (x_1^i(n), \dots, x_M^i(n))$ denote a decision vector (element of R^M) stored in the memory of decision maker i at time n . We also assume that each decision maker i stores in his memory another vector $(\lambda_1^i(n), \dots, \lambda_M^i(n))$ with his estimates of

$\frac{\partial J^1}{\partial x_i}, \dots, \frac{\partial J^M}{\partial x_i}$. Unlike the synchronous algorithm, we

do not require that a message be transmitted at each time stage and we allow communication delays. So let:

$p^{ki}(n)$ = the time that a message with a value of x_k was sent from processor k to

processor i , and this was the last such message received no later than time n .

$q^{ki}(n)$ = the time that a message with a value of

$\frac{\partial J^k}{\partial x_i}$ was sent from processor k to processor i , and this was the last such message received no later than time n .

For consistency of notation we let

$$p^{ii}(n) = q^{ii}(n) = n, \quad \forall i, \forall n. \quad (3.7)$$

With the above definitions, we have:

$$x_k^i(n) = x_k^k(p^{ki}(n)), \quad \forall n, \forall (k,i) \in E, \quad (3.8)$$

$$\lambda_i^k(n) = \frac{\partial J^k}{\partial x_i}(x^k(q^{ki}(n))), \quad \forall n, \forall (i,k) \in E. \quad (3.9)$$

in a distributed organization and may form the basis for a systematic approach to organizational design.

References

- Arrow, K.J., L. Hurwicz, (1960), "Decentralized and Computation in Resource Allocation," in Essays in Economics and Econometrics, R.W. Pfouts, Ed., Univ. of North Carolina Press, Chapel Hill, N.C. pp. 34-104.
- Bertsekas, D.P., (1982), "Distributed Dynamic Programming," IEEE Transactions on Automatic Control, Vol. AC-27, No. 3, pp. 610-616.
- Bertsekas, D.P., (1983), "Distributed Computation of Fixed Points," Math. Programming, Vol. 27, 1983, pp. 610-616.
- Gallager, R.G., (1977), "A Minimum Delay Routing Algorithm Using Distributed Computation," IEEE Transactions on Communications, Vol. COM-25, No. 1, pp. 73-85.
- Poljak, B.T., Y.Z. Tsypkin, (1973), "Pseudogradient Adaptation and Training Algorithms," Automatika i Telemekhanika, No. 3, pp. 45-68.
- Simon, H.A., (1980), The Sciences of the Artificial, MIT Press, Cambridge, MA.
- Tsitsiklis, J.N., (1983), "Problems in Distributed Decision Making," Technical Report, Laboratory for Information and Decision Systems, MIT, in preparation.
- Tsitsiklis, J.N., Bertsekas, D.P., Athanasiou, M., (1983), "Distributed Deterministic and Stochastic Optimization Algorithms with Applications in System Identification," Proceedings of the 1983 Moleccon Conference, Athens, Greece.

THE DESIGN OF INFORMATION STRUCTURES: BASIC ALLOCATION STRATEGIES FOR ORGANIZATIONS*

Debra A. Stabile

ICF, Incorporated
Washington, D.C.

Alexander H. Lewis

Massachusetts Institute of Technology
Cambridge, Massachusetts

ABSTRACT

The problem of designing the allocation of information processing tasks to organization members who interact with the organization's environment is formulated. Two information strategies are considered for reducing the load on each member while accomplishing the overall task: (a) creation of self contained tasks, and (b) creation of slack resources. The former leads to parallel processing, while the latter is accomplished through alternate processing rules. The two basic strategies can be integrated to produce a wide variety of information structures.

INTRODUCTION

In military organizations the ability to process information in an efficient (i.e., quick and accurate) manner can be of critical importance. Time constraints and limitations on the availability and capabilities of equipment and personnel may reduce the rate at which the organization's decision makers (DMs) can respond to information they receive. These constraints and limitations may force the DMs to become overloaded (i.e., each DM is assigned more tasks than he is able to execute in the prescribed time interval, while still maintaining a given performance level). Information reduction strategies may be employed to avoid overload. These strategies are implicit in the definition of information structures. A methodology is presented in this paper for designing the information structures for DMs who comprise the boundary between an organization and its environment. The objective of these designs is to process all information received by the organization efficiently, using the minimum number of organization members.

The design of organizations can be decomposed into two interrelated problems: the organizational form problem in which the information and decision structures are specified, and the organizational control problem in which the operating rules or procedures and the monitoring and enforcement strategies are determined. The manner in which information is received and processed by organization members is a key descriptor of the structure of an organization. Indeed, communication - the exchange of information and the transmission of meaning - is the very essence of an organization. To move from an unorganized state to an organized state requires the introduction of constraints and restrictions to reduce diffuse and random communications to channels appropriate for the accomplishment of organizational objectives (tasks) [1].

The organization, perceived as an open system [2], interacts with its environment; it receives signals or messages in various forms that contain information relevant to the organization's tasks. These messages must be identified, analyzed and transmitted to their appropriate destinations within

the organization. From this perspective, the organization acts as an information user.

How an organization receives signals from its environment has direct consequences on the internal structure of the organization and on its performance. The specific structure depends on the nature and characteristics of the signals that can be received, on the task to be performed, and on the capabilities and limitations of the individual members comprising the organization. By considering only the boundary between the organization and the environment, a major simplification occurs; the boundary itself can be thought of as a single echelon of organization members. While these members may occupy different positions in the internal organizational structure, the relevant characteristic is that they receive direct inputs from sources outside the organization. (Figure 1). In that sense, no member is subordinate to another; i.e., the members constitute a single echelon. However, individuals, or groups of individuals, can have very different capabilities and limitations that reflect, indirectly, their position and function in the organization. For example, they may be able to process only certain classes of signals (specialization) or deal with limited levels of uncertainty. Since it is important to remember that the single echelon may include commanders as well as operators of monitoring systems, executives as well as clerks, the term decisionmaker has been used to describe all members.




Figure 1 Environment and Single Echelon.
(Decisionmakers with hatched areas
comprise the single echelon.)

THE MODEL

The single echelon receives data from one or more sources external to the organization. Every δ_n units of time on the average, each source n generates symbols, signals, or messages x_{n1} from its associated alphabet X_n , with probability p_{n1} , i.e.,

$$p_{ni} = p(x_n = x_{ni}) ; x_{ni} \in X_n \quad i = 1, 2, \dots, \gamma_n \quad (1)$$

$$\sum_{i=1}^{\gamma_n} p_{ni} = 1 ; \quad n = 1, 2, \dots, N' \quad (2)$$

where γ_n is the dimension of X_n . Therefore, $1/\delta_n$ is the mean frequency of symbol generation from source n .

The task to be performed is defined as the processing of the input symbols x_n by the single echelon to produce output symbols. It is assumed that a specific complex task that must be performed can be modeled by N' such sources of data. Rather than considering these sources separately, one supersource \underline{x}' , composed of these N' sources, is created. The input symbol \underline{x}' may be represented by an N' -dimensional vector with each of the sources represented by a component of this vector, i.e.,

$$\underline{x}' = (x_1, x_2, \dots, x_{N'}) ; \underline{x}' \in X \quad (3)$$

To determine the probability that symbol \underline{x}' is generated, the independence between components must be considered. If all components are mutually independent, then $p_{\underline{x}'}$ is the product of the probabilities that each component of \underline{x}' takes on its respective value from its associated alphabet:

$$p_{\underline{x}'} = \prod_{n=1}^{N'} p_{ni} \quad (4)$$

When all components of the input vector are mutually independent, this is referred to as being of finest grain. In many situations, this assumption is unrealistic. It is more common to have some components probabilistically dependent.

If two or more components are probabilistically dependent on each other, but as a group are mutually independent from all other components of the input vector, then these dependent components can be treated as one supercomponent, with a new alphabet. Then a new input vector, \underline{x} , is defined, composed of the mutually independent components and these supercomponents. This new \underline{x} is on finest grain.

This model of the sources implies synchronization between the generation of the individual source elements so that they may, in fact, be treated as one input symbol. Specifically, it is assumed that the mean interarrival time for each component δ_n is equal to δ . It is also assumed that the generation of a particular input vector, \underline{x}_j , is independent of the symbols generated prior to or after it.

The DMs in the single echelon and the particular sequencing and allocation of information received by the DMs define the organizational form. The design of structures in which no DM is subordinate to any other DM will depend upon the constraints imposed by the organization. Two basic premises of the design method are that (a) in general, there is no best way to organize and (b) any way of organizing is not equally effective with respect to performing a specific task [3]. Characteristics of single echelon organizations

will be defined which suggest that certain structures will, in fact, be more effective than others for performing a specific complex task. For the types of organizations considered, the performance of a complex task is equivalent to the processing of information, where information is defined to be the data received by the DMs in the echelon. Galbraith has argued that variations in the amounts of information (data) that are not processed are primarily responsible for the variations in organizational forms [3]. Such variations are largely a result of the uncertainty associated with a given task. Uncertainty has been defined to be "the difference between the amount of data required to perform the task and the amount of data already possessed by the organization" [3].

It will be assumed that the number of DMs necessary in the echelon is greater than one. The underlying premise is that no DM alone is able to process the required amount of data while simultaneously achieving the required performance level.

Several factors affect the time required for a DM to process information. The uncertainty of the input symbol generated and the number of possible input-output responses are two of these factors. Keying in on these two factors, a processing time function which has an information interpretation can be introduced [4]. Let:

- s be the number of components the m -th DM processes
- $t^m(s)$ be a parameter which is a function of the number of components assigned to the m -th DM
- c^m be a constant
- $p(\underline{x}_{kj})$ be the probability associated with the j -th element of the k -th partition vector's alphabet.

Then:

$$\tau_{kj}^m = t^m(s) - c^m \log p(\underline{x}_{kj}) \quad (5)$$

If this processing time function is averaged over all possible elements of the input symbol to be processed [4], then

$$\begin{aligned} \tau_k^m &= \sum_j p(\underline{x}_{kj}) \tau_{kj}^m \\ &= t^m(s) - c^m \sum_j p(\underline{x}_{kj}) \log p(\underline{x}_{kj}) \end{aligned} \quad (6)$$

In information theory, the last term of equation (6), is defined to be the entropy H associated with the group of components, k . It follows that:

$$\tau_k^m = t^m(s) + c^m H_k \quad (7)$$

The quantity t^m is a monotonic nondecreasing function of the number of components, s , assigned to the m -th DM. It describes the fixed cost in terms of time required to process any component regardless of the amount of information contained in it. A multiplicative function of the number of components

assigned to the m -th DM is often assumed, i.e.,

$$t^m(a) = t_0^m \quad (8)$$

where t_0^m is a constant.

The parameter c in (7) is assumed to be a constant for each DM considered. To compute the average processing time for DM _{m} to process the group of components, k , the entropy H_k must first be computed. Many distinct groupings of the components of a can be constructed and an entropy must be computed for each of those groupings. For each supercomponent, $x(n)$, the corresponding entropy is:

$$H_n = - \sum_{i=1}^r p(x(n) = x_i(n)) \log(p(x(n) = x_i(n))) \quad (9)$$

Since it has been assumed that all components in the input vector are mutually independent, the entropy of any set of components L_k is equal to the sum of the entropies of each of the components in the set, i.e.,

$$H_{L_k} = \sum_{n \in L_k} H_n \quad (10)$$

This results in substantial savings in computation. Rather than needing to compute the probabilities of each element of each alphabet for each distinct grouping of components in order to compute the associated entropy, it is only necessary to compute the entropy for each component of the input vector.

OVERLOAD

When a DM is overloaded, he can react in one of several ways. He may decide to reduce the amount of data he has to process by either randomly (rejection) or selectively (filtering) omitting data. The amount of data he may be required to process may be reduced also by having it preprocessed. He may decide to reduce the number of categories of discrimination (i.e., approximate the inputs) or he may reduce the required level of accuracy for processing the data and, in so doing, reduce the number of different outputs. If these alternatives seem unsatisfactory, he may decide to receive all the data allowing queues to build up, delaying the processing during periods of peak loads and attempting to catch up during time intervals when input symbols are assigned to other DMs. Otherwise, the DM may simply choose not to perform the task [3].

An alternative to having the data preprocessed is the employment of multiple parallel channels [2], [6]. This paper will focus on the use of parallel channels (DMs) which guarantee all the data received are processed immediately. Queuing also guarantees all the information is processed, but not immediately. The other approaches to reducing overload are unappealing to the organization designer as a first choice in two ways: (1) data are omitted and (2) wasteful expenditures are incurred since the value of the data sent but not used is lost.

The concept of parallel DMs is analogous to the idea of distributed information processing with each DM performing a subtask. Thus the idea of considering a single echelon structure only is a relevant issue for any organization. Many studies have revealed that

as the uncertainty of the tasks increases, the "flatter" (i.e., more distributed) an organization should become with respect to its DMs [2].

Galbraith [3] has suggested two information reduction strategies for organizations to address this issue: (1) Creation of Self-Contained Tasks, and (2) Creation of Slack Resources. In the first strategy, the original task is divided into a set of subtasks. This reduces (a) the diversity of outputs each DM in the organization can produce and (b) the number of different inputs, since the DM need only receive the inputs pertaining to the given subtask. The slack resource to be considered for the second strategy will be time and will involve a reduction in the required level of performance with respect to time. For example, if it was originally required to have the task completed in time δ , this time may be extended to some multiple of δ .

Two types of processing modes will be considered, parallel processing, which is associated with the first strategy, and alternate processing, which is associated with the second strategy. These two modes are the fundamental strategies employed to guarantee all data are processed without overloading any DM in the echelon. The two fundamental modes can be integrated in various ways so as to develop more complex organizational forms.

PARALLEL PROCESSING

In a parallel processing structure, partitions of the input symbol are selected and assigned to the DMs in the organization. The group of DMs who, together, process the entire input vector form the single echelon. Each DM is constrained to process a partition of components from those that do not overload him (i.e., from those that result in a mean processing time of δ or less).

Mathematical programming in an appropriate modeling approach for this class of problems. This approach seeks "the optimum allocation of limited resources among competing activities under a set of constraints imposed by the nature of the problem being studied" [5]. In this context, the components of the input vector correspond to the limited resources, the DMs correspond to the competing activities and the constraint sets include processing time capabilities and specialization limitations of each DM.

Individual components are selected and assigned to each DM in the model. The group of components assigned to the m -th DM defines the partition vector. The conditions for selecting and assigning components are:

- (a) every component is processed, and
- (b) no DM is overloaded.

Let Y_{nm} be a binary variable which equals one if the n -th component is assigned to the m -th DM, zero otherwise. To guarantee that every component is processed once and only once, the following set of constraints is established:

$$\sum_{m=1}^M Y_{nm} = 1 \quad ; \quad n = 1, 2, \dots, N \quad (11)$$

where

$$Y_{nm} = 0,1 \quad (12)$$

A network structure which links every component to every DM is shown in Figure 2.




Figure 2 Parallel Processing Structure.

A DM will not be overloaded if the average time he requires to process the components assigned to him does not exceed δ . The mean processing time is given by (eqs. (7) and (8))

$$\bar{\tau}^m = s t^m + c^m H_n \quad (13)$$

where s is the number of components assigned to the DM. Since the components have been assumed to be mutually independent, the entropy H is equal to the sum of the entropies of the s components. Since the components assigned to the m -th DM are not known a priori, a binary indicator variable Y_{nm} is introduced which includes the time for processing component x_n only if it is assigned to the DM:

$$(t^m + c^m H_n) Y_{nm} = \begin{cases} (t^m + c^m H_n) & \text{if } Y_{nm} = 1 \\ 0 & \text{if } Y_{nm} = 0 \end{cases} \quad (14)$$

Furthermore:

$$\bar{\tau}^m = \sum_{n=1}^N (t^m + c^m H_n) Y_{nm} \leq \delta ; m = 1, 2, \dots, M \quad (15)$$

i.e., $\bar{\tau}^m$ must be less than or equal to δ to guarantee the m -th DM is not overloaded.

The objective function for this problem is to minimize the number of DMs required to process all the components without overload. The information structure can be constructed from the optimal solution to this problem. The optimal solution would identify who and how many DMs are included in the echelon and what subtasks they are performing.

ALTERNATE PROCESSING

Information structures based on alternate processing are appropriate when the input vector cannot be partitioned (i.e., when the strategy of creating self-contained tasks cannot be used to avoid overload). This strategy involves the creation of a slack resource, in this case, time. Thus, each DM is given more time to process the input assigned to him, which introduces a delay strictly greater than δ . Figure 3 illustrates alternate processing.

A deterministic strategy is one in which the ordering of the assignment of the input vectors to the DMs is fixed. In order to specify the optimal information structures associated with this strategy, it is necessary to determine simultaneously:

- the minimum number of DMs, M^* , necessary to process the input vectors without any one being overloaded, and
- the frequency, q_m , with which each of these DMs receives an input vector.




Figure 3 Alternate Processing Structure.

A very simple method for solving this problem exists. The overload constraint requires:

$$q_m \leq \delta / \bar{\tau}^m \quad (16)$$

where $\bar{\tau}^m$ is the average time for the m -th DM to process an input vector. Without any loss of generality, the DMs may be re-indexed according to their processing time functions (i.e., let the first DM be the most efficient and the m -th DM be the least efficient, so that $\bar{\tau}^1 \leq \bar{\tau}^2 \leq \dots \leq \bar{\tau}^M$). The other constraint on the problem is that all of the data be processed:

$$\sum_{m=1}^{M^*} q_m = 1 \quad (17)$$

where M^* has yet to be determined. The solution proceeds by choosing DMs in order of efficiency until:

$$\sum_{m=1}^{\lambda+1} \delta / \bar{\tau}^m > 1 \geq \sum_{m=1}^{\lambda} \delta / \bar{\tau}^m \quad (18)$$

If the right hand side of equation (18) is an equality, then the minimum number of DMs, M^* , necessary to process the input vectors without overload is equal to λ , and

$$q_m = \begin{cases} \delta / \bar{\tau}^m & \text{if } 1 \leq m \leq \lambda \\ 0 & \text{otherwise} \end{cases} \quad (19)$$

If the right hand side of equation (18) is a strict inequality, then the minimum number of DMs, M^* ,

equal to $\lambda+1$. Because

$$\sum_{m=1}^{\lambda+1} \delta/\tau^m > 1,$$

q_m must be defined as

$$q_m = \begin{cases} \delta/\tau^m - \rho^m & \text{if } 1 \leq m \leq \lambda+1 \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

are

$$\rho^m = \rho/\lambda+1 \quad m = 1, 2, \dots, \lambda+1$$

and

$$\rho = \left(\sum_{m=1}^{\lambda+1} \delta/\tau^m \right) - 1 \quad (21)$$

These ρ^m may be set to ensure that all of the q_m are rational, so that a cyclical strategy can be used.

A cyclical strategy is defined as a strategy in which the ordering of the assignment of the input vectors to the DMs is repeated every α' input vectors. In the case that the right hand side of (16) is a quality and at least one q_m is irrational, a cyclical strategy cannot be used; but it may be argued that since the τ^m are usually estimated rather than precisely calculated, a cyclical strategy can always be used. In this case, α' , the number of inputs in one cycle, is the lowest common denominator of the q_m 's. The information structures for a deterministic cyclical strategy may now be completely specified.

Define F to be the ordered set of indices on one cycle of α' input vectors; that is,

$$F = \{f \mid f = 1, 2, \dots, \alpha'\} \quad (22)$$

Now let F^m be a subset of F where

$$F^m = \{f \in F \mid \text{input } x_f \text{ is assigned to DM } m\} \quad m = 1, 2, \dots, M^* \quad (23)$$

With the indicator ϕ_f^m defined as:

$$\phi_f^m = \begin{cases} 1 & \text{if } f \in F^m \\ 0 & \text{otherwise} \end{cases} \quad m = 1, 2, \dots, M^* \quad (24)$$

the only requirement on the assignment of input to the m -th DM is that:

$$\sum_{f=1}^{\alpha'} \frac{\phi_f^m}{\alpha'} = q_m \quad m = 1, 2, \dots, M^* \quad (25)$$

Since input vectors only arrive once every δ time

units, x_f is assigned at time $t = (k\alpha' + f)\delta$, where $k=0, 1, \dots$ determines the number of cycles that have been completed at time t . The optimal solution identifies the number and properties of the many DMs who are included in the single echelon and the frequency with which each DM receives an input symbol.

DESIGN METHODOLOGY

The modeling and analysis of the information structures of single echelon organizations lead to a five step design methodology. The first step is the analysis of the properties of the inputs to the organization, i.e., of the task to be performed, and the mathematical modeling of the source so that the components of the symbol vector are mutually independent. The second step consists of the selection of the appropriate information reduction strategies in view of the DMs that are available to the designer. The basic ones considered are parallel and alternate processing. Integrated strategies that consist of various combinations of alternate and parallel processing are also considered. It is here that the organization designer's understanding of the task to be performed by the organization is crucial. The third step consists of the formulation of the mathematical model that represents the (integrated) information reduction strategy selected in the second step. All constraints that characterize the task, the organization, and the individual DMs are also expressed analytically. The fourth step is the posing of the optimization problem. Given the mathematical formulation of the single echelon model, an objective function must be defined. While in specific cases, solutions to the optimization problem can be obtained from straightforward computations, the general formulation lends itself to mathematical programming techniques. In particular, generalized network (GN) formulations proved most attractive because of the efficient algorithms which exist to solve them [7]. Knapsack problems and mixed integer linear programs can also be considered. The fifth step consists of expressing the solution to the optimization problem in the form of an information structure. The structure is then evaluated to determine whether the trade-offs between the number of DMs used and the delays are acceptable; if not, then the designer should return to step 2 and revise the proposed information reduction strategy.

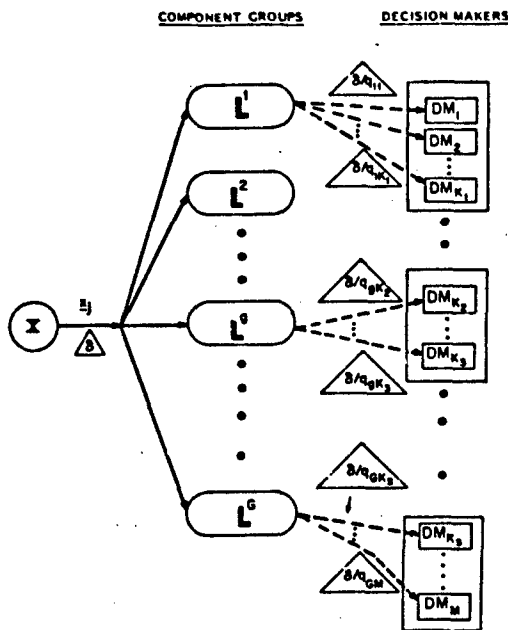
In the next section, an example is presented that illustrates the application of the five step methodology.

EXAMPLE

Consider G distinct sources, each source generating a vector of signals. The task is such that each source output has to be processed intact (i.e., it cannot be partitioned). There are M DMs who can receive the generated signals and none of these DMs can process the output of any of the G sources without being overloaded. A parallel/alternate information structure seems appropriate. A generic form of this structure is shown in Figure 4.

Step 1: TASK MODELING

The supercomponent consists of G synchronized sources that generate vector signals. The mean signal generation rate is δ^{-1} . The elements of the input vector can be partitioned into G sets, each set corresponding to the output of each of the individual sources. This is the finest grain decomposition of the input.



where Δ input symbol generation time

Figure 4: Parallel/Alternate Information Structure.

Step 2: INFORMATION REDUCTION STRATEGY

It is assumed that up to M distinct decisionmakers may be used in the single echelon.

The decomposition of the input vector allows for the parallel processing of the signals generated by the G sources. No further division into subtasks is possible. Since every one of the G subtasks arriving at a rate δ^{-1} cannot be processed by any one DM without causing overload, the second information reduction strategy (creation of slack resources) must be used. Alternate processing of signals generated by each source will allow additional time for each DM to do the processing and, therefore, overload may be avoided. The resulting processing mode is an integrated parallel/alternate processing.

Step 3: MATHEMATICAL MODEL

- (a) Since alternate processing is assumed for the output of each source, the requirements that all signals be processed reduces to the condition that the sum of the symbol assignment frequencies for the output of each source, q_{gm} , must be equal to unity, i.e.,

$$\sum_{m=1}^M q_{gm} = 1 \quad g = 1, 2, \dots, G \quad (26)$$

- (b) In order that no DM be overloaded, the frequency with which each DM receives a signal for processing should be sufficiently low so that his mean processing time does not exceed the effective mean interarrival time. This condition becomes:

$$0 \leq q_{gm} \leq \delta / \tau_m^g \quad g = 1, 2, \dots, G \quad (27)$$

$$m = 1, 2, \dots, M$$

- (c) Any DMs that are assigned input for processing with zero frequency are excluded from the single echelon. Furthermore, each DM is allowed to receive inputs from at most one of the G sources. Constraints (28) and (29) guarantee these conditions, where the binary variable Y is zero when the m-th DM is assigned the output of the g-th source.

$$\sum_{g=1}^G Y_{gm} = G - 1 \quad m = 1, 2, \dots, M \quad (28)$$

$$Y_{gm} \cdot q_{gm} = 0 \quad g = 1, 2, \dots, G; \quad m = 1, 2, \dots, M \quad (29)$$

Note that while the inputs from the sources are received by the single echelon simultaneously, the outputs are not synchronized. Indeed, each DM introduces a different delay; the maximum delay is given by the maximum value of

$$\delta (1 - q_{gm}) / q_{gm} \quad (30)$$

over all m. If this delay is unacceptable, then more efficient DMs are needed.

Step 4: OPTIMIZATION

In this problem, where M distinct DMs are available, the number of decisionmakers that can process the incoming signals is a reasonable objective function to be minimized. Since the DMs do not have identical properties (parameters), the problem cannot be decoupled into G distinct optimization problems, even though no decisionmaker is allowed to process signals from more than one source. The resulting mathematical programming (MP) problem must be solved algorithmically. Furthermore, the nonlinearity of one of the constraints makes it a difficult one to solve.

Step 5: INFORMATION STRUCTURE

The solution to the nonlinear MP, namely, the values of M^* , q_{gm}^* , and Y_{gm}^* , defines the information structure. The single echelon is composed of only those DMs for which the corresponding frequency q_{gm} is strictly positive. The information structure, Figure 4, specifies the decisionmaker, the group components L^g he processes, and the frequency q_{gm} with which he is assigned these inputs.

CONCLUSION

An approach to the design of information structures for a single echelon organization has been presented. This approach is based on the properties of the inputs, the characteristics of the available DMs, and the constraints imposed on the organization by the task. Two basic information reduction

strategies, creation of self-contained tasks and creation of slack resources, were modeled as parallel and alternate processing, respectively. It was then shown that complex information structures can be constructed using combinations of parallel and alternate processing. The former is appropriate when an overall task can be divided into subtasks; the latter, when delays in producing an output can be tolerated and the task cannot be divided.

The next major step in this research is the integration of the single echelon with other parts of the organization. The single echelon is responsible for transmitting the processed inputs to the appropriate destinations within the organization. This transmission of processed data to other members in the organization is referred to as serial processing.

The design of multiechelon structures requires each echelon to process its information without overload. The constraints on each echelon, however, must be inferred from the constraints that are imposed on the overall organization. This introduces a higher level of complexity to the design problem.

ACKNOWLEDGMENT

This work was supported by the Office of Naval Research under contracts ONR/N00014-77-C-0532 and ONR/N00014-83-K-0185 and the Air Force Office of Scientific Research under contract AFOSR-80-0229.

- [1] D. Katz and R. L. Kahn, The Social Psychology of Organizations, John Wiley and Sons, New York, 1966.
- [2] E. E. Lawler, III and J. G. Rhode, Information and Control in Organizations, Goodyear Publishing Company, Inc., Pacific Palisades, California, 1976.
- [3] J. Galbraith, Designing Complex Organizations, Addison-Wesley, New York, 1973.
- [4] T. B. Sheridan and W. R. Ferrell, Man-Machine Systems, The M.I.T. Press, Cambridge, MA, 1974.
- [5] R. V. Hogg and A. T. Craig, Introduction to Mathematical Statistics, Macmillan, New York, 1968.
- [6] R. F. Drenick, "Organization and Control," in Directions in Large Scale Systems, Y. C. Ho and S. K. Mitter, Eds., Plenum Press, New York, 1976.
- [7] D. A. Stabile, "The Design of Information Structures: Basic Allocation Strategies for Organizations," M.S. Thesis. Report LIDS-TH-1098, Laboratory for Information and Decision Systems, M.I.T., Cambridge, MA, June 1981.

AD P002910

INFORMATIONAL ASPECTS OF A CLASS OF SUBJECTIVE GAMES OF INCOMPLETE INFORMATION*

Dr. Demosthenis Teneketzis

Dr. David A. Castanon

ALPHATECH, Inc.
2 Burlington Executive Center
111 Middlesex Turnpike
Burlington, Massachusetts 01803

1. INTRODUCTION

The key assumptions upon which the development of game theory was initially based are the following:

- A1) The rules of the game are common knowledge to all the players of the game.
- A2) The players have the same perception (model) of the game.
- A3) Players are fully committed to a priori strategies.
- A4) Players are rational.

Assumptions (A1)-(A3) are quite restrictive as they do not hold in many real-life economic, political, military and other social situations. This is why, as game theory developed, attempts were made to relax some of these assumptions. Assumption A3 was a consequence of the normalization principle of Von-Neumann, [12, pp 79-84] which roughly says that given an extensive game, one can always reduce it to an equivalent game in normal form involving only strategies and payoffs and where all dynamic and informational aspects of the original problem have been expressed in the form of strategies by considering all the possible actions of all the players under all possible circumstances. Aumann and Maschler [1] were the first to point out via a simple counterexample the inappropriateness of the normalization principle under certain conditions; since then considerable developments followed by relaxing the assumption of prior commitment [2]-[6]. Harsanyi [7] and Aumann-Maschler et al. [8] pointed out that in some military problems, players may lack full information about the payoff functions of other players, or about the physical facilities and strategies of other players, or even about the amount of information that other players have about the various aspects of the game situation. Thus, Harsanyi [7] first relaxed assumption A1 and formulated and developed models of games of incomplete information. Considerable progress has been achieved in the theory of games of incomplete information since Harsanyi's original formulation (see [8]-[10] and references therein).

In this paper, we shall relax assumptions A1 and A2, we shall modify assumption A4, and we shall formulate a class of games which we shall call "Subjective Games". We shall consider that the players have different perceptions (models) of the game that is being played. We shall further assume that each player considers that the other players have models of the

game which is the same as his own model. Moreover, we shall consider that each player is rational within his own subjective view of the decision problem. Under these assumptions, we shall study one-stage games as well as repeated games. Various interesting issues arise because of the new assumptions:

- Q1. How are equilibrium strategies defined for subjective games?
- Q2. How do these equilibrium strategies relate to the equilibrium strategies of the games studied so far?
- Q3. Do players realize during the play of the game that they have different models?
- Q4. If, during the play of the game, the players realize that they have different models, how do they modify their models and their strategies?
- Q5. Does repetition of the game result in cooperation as in the case of the games studied so far (e.g., [11])? And does repetition of the game alleviate differences in the models of the players and allow players to agree on an equilibrium strategy?
- Q6. How does bargaining take place in subjective games? What is the effect of the bargaining model used on the outcome of bargaining?
- Q7. Is it possible to characterize the set of all equilibria for repeated subjective games?

To study and understand some of these questions, we shall consider a special class of games, namely 2×2 , two person non-zero sum games of incomplete information where the payoff matrices have a special structure.

The rest of the paper is organized as follows: In Section 2, we present the model for subjective games; specifically, we state the basic assumptions under which the theory will be developed. In Section 3, we briefly discuss games of incomplete information and point out the differences between Harsanyi's model and our model. In Section 4, we consider static subjective non-cooperative games of incomplete information; we study informational aspects of these games and contrast the results to those of classical games. Conclusions are presented in Section 5. This paper is part of a larger report [16] where infinitely repeated non-cooperative subjective games of incomplete information are studied and solved. The section on repeated subjective games has been omitted due to space limitations

*Work sponsored by Office of Naval Research, Department of the Navy, 500 N. Quincy Street, Arlington, VA 22217 Contract No. N000014-82-C-0693.

The contributions of this paper are as follows:

(i) It presents a model for subjective games which relaxes some of the restrictive key assumptions of game theory previously made.

(ii) It studies a simple class of static subjective games and their informational aspects; it shows the effect of the subjective models on the solution of the game and contrasts the value of information in these games to the value of information in classical static Bayesian games.

2. THE MODEL FOR SUBJECTIVE GAMES

We shall develop our theory of subjective games based on the following key assumptions:

- S1. Players have different perceptions (different models) of the game.
- S2. Each player thinks that the other players' perception (model) of the game is the same as his.
- S3. Players are Bayesian.
- S4. Each player is rational within his own subjective view of the game.

Assumptions S1 and S3 imply that each player j will assign a subjective joint probability distribution P_j to all variables unknown to him. A subjective probability distribution P_j entertained by player j is defined in terms of his own choice behavior and may be considered as his personal estimate of the variables unknown to him [13]. In contrast to S1 and S3, assumption A2 combined with Bayesian players implies the existence of an objective probability distribution P^* which is defined in terms of the long-run frequencies of the relevant events; these frequencies are established by an independent observer, the umpire of the game.

Assumption S2 implies that the rules of the game are not common knowledge to all the players, since each player thinks that the other players' perception of the game is the same as his, yet this may not be true. Assumptions S1 and S2 were previously used in the context of distributed estimation and detection [14]. Some of the interesting issues that arise because of S1 and S2 are the following:

- I1. How do players interpret moves of other players or messages received from other players?
- I2. Do players realize during the play of the game that they have different models or do they play the game and achieve an equilibrium which is acceptable within the terms of their own models?
- I3. If players discover that they have different models, how do they reconcile their differences?

Some of the above issues were studied and answered in [14] in the context of distributed estimation and detection.

Thus, comparing S1-S4 with A1-A4, we see that only the rationality assumption holds. Under assumptions S1-S4, several interesting issues like Q1-Q7 discussed in Section 1 arise. We shall study some of these issues in the context of a specific class of games, namely games of incomplete information.

3. GAMES OF INCOMPLETE INFORMATION

Games of Incomplete Information (GII) were first formulated by Harsanyi [7]. In his original formulation, Harsanyi called GII those situations where "the participants lack full information about some important aspects of the game they are playing". Harsanyi considers a n -player game and makes the following key assumptions in his formulation:

- H1. Each player has certain information about the game. This information is described by a vector

$$c_i = (a_i^T, b_i^T)^T.$$

The component a_i represents the private information of player i about the payoff functions J_1, J_2, \dots, J_n of the n players (the component b_i will be defined below).

- H2. In dealing with incomplete information, each player takes a Bayesian approach. That is, each player assigns a subjective probability distribution P_i .

$$P_i = P_i(c_1, c_2, \dots, c_{i-1}, c_{i+1}, \dots, c_n)$$

$$\Delta P(c^i) \quad (3-1)$$

to all the variables unknown to him and attempts to maximize the mathematical expectation of his own payoff J_i in terms of this probability distribution P_i .

- H3. The other players do not know the subjective probability P_i used by player i . But each player j ($j \neq i$) can write P_i in the form

$$P_i(c^i) = R_i(c^i | b_i) \quad (3-2)$$

where R_i is a function whose mathematical form is known to all players, and b_i is a vector consisting of those parameters of the function P_i that are known only to player i . Thus, each player has a probability distribution over the subjective distributions P_i of the other players, and these probability distributions are common knowledge to all the players since the mathematical form of each R_i is known to all players.

Harsanyi shows that under H1-H3, the GII is Bayes equivalent to a Game of Complete Information (GCI) for each player, (i.e., completely equivalent to a GCI for each player from a game theoretic standpoint), if there is a probability distribution $R^*(c_1, c_2, \dots, c_n)$ such that

$$R^*(c^i | a_i, b_i) = R_i(c^i | a_i, b_i) V_i \quad (3-3)$$

Furthermore he shows that the normal form for the equivalent Bayesian game is in many cases an unsatisfactory representation of the game situation and has to be replaced by other representations, e.g., the semi-normal (or extensive) form, where Bayesian games must be interpreted as games with delayed commitment.

Thus, Harsanyi's formulation drastically departs from the standard formulation of game problems as it relaxes two of the basic assumptions, namely A2 and A3.

Following Harsanyi's original formulation, GII were analyzed by many researchers (see [9] - [10] and references therein), and the sets of equilibria for zero and non-zero sum, static and repeated GII were determined.

The study of subjective games under assumptions S1-S4 will be based on subjective GII. Our model differs from that of Harsanyi in one basic point. Harsanyi's formulation assumes that the rules of the game are common knowledge to all the players (i.e., assumption A1 holds). This happens because each player knows what information is available to other players and he also knows the functional form of the subjective joint probability distribution entertained by each other player. Our formulation assumes that the rules of the games are not common knowledge to all players. This happens because each player has his own subjective probability about the game and in addition he thinks that all the other players share the same view about the game as he does. In spite of this difference a lot of results from GII will be useful in analyzing subjective GII and in showing similarities and differences between these two classes of games.

4. STATIC SUBJECTIVE GAMES OF INCOMPLETE INFORMATION

In this section, we consider static non-zero sum games of incomplete information. We present a class of subjective games with simple solutions, we study the value of public, private and secret information, and we contrast the results with those of the classical static Bayesian games.

4.1 Problem Formulation

We consider the following static two person non-zero sum game. Chance selects one of two games with the following payoff matrices:

Game 1

$$\begin{array}{cc} & \sigma & \tau \\ \lambda & (c_{11}, c_{11}) & (c_{21}, c_{12}) \\ \mu & (c_{12}, c_{21}) & (c_{22}, c_{22}) \end{array} \quad (4-1)$$

Game 2

$$\begin{array}{cc} & \sigma & \tau \\ \lambda & (c_{12}, c_{12}) & (c_{22}, c_{11}) \\ \mu & (c_{11}, c_{22}) & (c_{21}, c_{21}) \end{array} \quad (4-2)$$

We further assume that

$$c_{11} > c_{21} > c_{12} > c_{22} \quad (4-3)$$

$$c_{12} + c_{21} > c_{11} + c_{22} \quad (4-4)$$

Player 1 can choose action (λ, μ) and player 2 can choose action σ, τ .

Note that because of 4-3 each player has a dominant strategy, in each one of the two games. So far, the statement of the problem and the assumptions (4-3) - (4-4) are essentially the same as in [15]. However, contrary to [15], we now assume that the two players have a different perception of the game, namely:

- B1. Player 1 thinks that chance selects game 1 with probability

$$P(1) = p > 1/2 \quad (4-5)$$

- B2. Player 2 thinks that chance selects game 1 with probability

$$P(1) = q < 1/2 \quad (4-6)$$

- B3. Each player thinks that the other player's model of the game is the same as his own model

- B4. Chance actually selects game 1 with probability $P(1) = r$

We further consider that

- B5. Each player is rational within his own subjective view of the game.

Assumptions B1-B5 are the analogs of S1-S4 for the specific game. Under these assumptions, we shall study the solution of the game as well as various informational aspects of it.

4.2 The Solution of the Game

We study the problem formulated above under three different types of information that a player may receive:

1. Public Information

In this case, both players are informed about the outcome of the chance move.

2. Private Information

In this case, one player (say player 1) is informed about the outcome of the chance move whereas the other (player 2) is not. Moreover, the uninformed player knows that his opponent is informed.

3. Secret Information

In this case, one player is informed about the outcome of the chance move whereas the other is uninformed. Moreover, the uninformed player is unaware that his opponent is informed.

For each of these games, we present the solution of the game. Moreover, in order to compute the value of information for each case, we need to find the solution of the game where no player is informed. Thus, we solve the game for this case too. For each case, each player perceives a game which is different from the game perceived by the other player. Therefore, in some cases, we shall have to consider two different games in order to find the solution for each case.

1. Public Information

In this case, player 1 plays λ in Game 1 and μ in Game 2. Player 2 plays σ in Game 1 and τ in Game 2. Thus, the outcomes are (c_{11}, c_{11}) in Game 1 and (c_{21}, c_{21}) in Game 2 and the payoff of the players is

$$J_1^B = J_2^B = rc_{11} + (1-r)c_{21} \quad (4-7)$$

2. Private Information

2a. Assume at first that player 1 is informed. Then he plays σ in Game 1 and μ in Game 2. Player 2 has two options. Choose either σ or τ . If he chooses σ , he expects a payoff equal to

$$q c_{11} + (1-q) c_{22} \quad (4-8)$$

If he chooses τ then he expects a payoff equal to

$$q c_{12} + (1-q) c_{21} \quad (4-9)$$

Because of (4-4) and (4-6), player 2 prefers to play τ . Thus, the payoffs of the two players are

$$J_1^P = c_{21} \quad (4-10)$$

$$J_2^P = r c_{12} + (1-r) c_{21} \quad (4-11)$$

2b. If player 2 is informed, then he plays σ in Game 1 and τ in Game 2. An argument similar to that of Case 2a shows that if

$$p < \frac{c_{11} - c_{22}}{c_{11} + c_{21} - c_{12} - c_{22}} \quad (4-12)$$

then player 1 will play μ . Otherwise, he will play λ . The expected payoffs for player 1 are then:

$$J_1^P = r c_{12} + (1-r) c_{21} \quad (4-13)$$

and

$$J_1^P = r c_{11} + (1-r) c_{22} \quad (4-14)$$

respectively.

The payoff for player 2 is:

$$J_2^P = c_{21} \text{ (corresponding to } \mu) \quad (4-15)$$

and

$$J_2^P = c_{11} \text{ (corresponding to } \lambda) \quad (4-16)$$

3. Secret Information

3a. Assume at first that player 1 is secretly informed about the outcome of the chance move. Then, because in each game he has a dominant strategy, he plays λ in Game 1 and μ in Game 2. Player 2 thinks that the following game is being played:

$$\begin{bmatrix} (q c_{11} + (1-q) c_{12}), (q c_{11} + (1-q) c_{12}) \\ (q c_{21} + (1-q) c_{22}), (q c_{12} + (1-q) c_{11}) \\ (q c_{12} + (1-q) c_{11}), (q c_{21} + (1-q) c_{22}) \\ (q c_{22} + (1-q) c_{21}), (q c_{22} + (1-q) c_{21}) \end{bmatrix} \quad (4-17)$$

Because of (4-3), (4-6) each player has a dominant strategy in this game. Thus, player 2 plays τ . The payoffs of the two players are:

$$J_1^C = c_{21} \quad (4-18)$$

$$J_2^C = r c_{12} + (1-r) c_{21} \quad (4-19)$$

3b. If player 2 is secretly informed then, because in each game he has a dominant strategy, he plays σ in Game 1 and τ in Game 2. Player 1 acts considering that the following game is being played:

$$\begin{bmatrix} (p c_{11} + (1-p) c_{12}), (p c_{11} + (1-p) c_{12}) \\ (p c_{21} + (1-p) c_{22}), (p c_{12} + (1-p) c_{11}) \\ (p c_{12} + (1-p) c_{11}), (p c_{21} + (1-p) c_{22}) \\ (p c_{22} + (1-p) c_{21}), (p c_{22} + (1-p) c_{21}) \end{bmatrix} \quad (4-20)$$

Because of (4-3), (4-5), λ is a dominant strategy. Thus, player 1 plays λ . The payoffs of the two players are:

$$J_1^C = r c_{11} + (1-r) c_{22} \quad (4-21)$$

$$J_2^C = c_{11} \quad (4-22)$$

Finally, in order to compute the value of information for each case, we need to solve the game for the case where none of the players is informed about the outcome of the chance move. In this case, player 1 thinks that the game whose payoff matrix is given by (4-20) is being played, whereas player 2 thinks that the game whose payoff matrix is given by (4-17) is being played. Therefore, player 1 plays λ and player 2 plays τ . The payoffs of the two players are:

$$J_1^0 = r c_{21} + (1-r) c_{22} \quad (4-23)$$

and

$$J_2^0 = r c_{12} + (1-r) c_{11} \quad (4-24)$$

Let us discuss now some interesting features of the solutions of these games. At first note that each payoff bimatrix is symmetric, hence in each one of the two games, the players are interchangeable. Thus, one expects that for the classical Bayesian game, in the case of public or secret information, the behavior of the informed and the uninformed player will be independent of who is the informed and who is the uninformed player. For example, in the case of private or secret information if player 1 were the uninformed player and played λ (first row), we would expect that if the situation were reversed and player 2 became the uninformed player he would play σ (first column). Also, in the case where no player was informed about the outcome of the chance move, the dominant strategies would be (λ, σ) or (μ, τ) , hence the outcome of the game would be either (λ, σ) or (μ, τ) . Consequently, the value of private, secret or public information would be the same for both players. It can be easily checked that this is indeed the case when $p=q=r$. However, this behavior is not observed when each player has his own subjective model of the game. When player 1 is privately informed about the chance move, player 2 always chooses τ (second column); on the other hand, if player 2 is privately informed about the outcome of the chance move, player 1 does not always play μ (second row). When player 1 is the secretly informed player, player 2 always plays τ (second column); if player 2 is the

secretly informed player, player 1 always plays λ (first row). When no player is informed about the outcome of the chance move, the outcome of the game is (λ, r) . These facts indicate that the value of private and secret information is now different for each player. (Obviously, the asymmetry is induced by our assumption $0 < q < 1/2 < p < 1$.) Indeed, this is true as we shall see in the next section where we study the value of the information for the above game.

4.3 The Value of Information

The value of information is defined in general as follows:

V_1 = (Payoff of player 1 when he knows the outcome of the chance move) - (Payoff of player 1 when no player is informed about the outcome of the chance move.)

We shall compute the value of public, private and secret information for each player.

1. Value of public information.

Because of (4-7), (4-23) and (4-24) we find that

$$V_1^B = r(c_{11} - c_{21}) + (1-r)(c_{21} - c_{22}) \quad (4-25)$$

$$V_2^B = (2r-1)c_{11} + (1-r)c_{21} - rc_{12} \quad (4-26)$$

2. Value of Private Information

Because of (4-10), (4-14) - (4-16), (4-23) and (4-24) we find that:

$$V_1^P = (1-r)(c_{21} - c_{22}) \quad (4-27)$$

$$V_2^P = c_{21} - rc_{12} - (1-r)c_{11}$$

$$\text{if } P < \frac{c_{11} - c_{22}}{c_{11} + c_{21} - c_{12} - c_{22}} \quad (4-28)$$

$$V_2^P = r(c_{11} - c_{12})$$

$$\text{if } P > \frac{c_{11} - c_{22}}{c_{11} + c_{21} - c_{22} - c_{12}} \quad (4-29)$$

3. Value of Secret Information

Because of (4-18), (4-22), (4-23) and (4-24), we get:

$$V_1^S = (1-r)(c_{21} - c_{22}) \quad (4-30)$$

$$V_2^S = r(c_{11} - c_{12}) \quad (4-31)$$

Thus, for the class of games considered in this section, the value of public, private and secret information differs from player to player, whereas in the classical Bayesian framework the value of public, private and secret information does not depend on who is the informed and who is the uninformed player.

Consider the case where player 2 is privately informed,

$$P < \frac{c_{11} - c_{22}}{c_{11} + c_{21} - c_{12} - c_{22}}$$

and

$$r = \frac{1}{2} \quad (4-32)$$

Then, the value of information for player 2 is, according to (4-28),

$$V_2^P = c_{21} - \frac{1}{2}c_{12} = \frac{1}{2}c_{11} \quad (4-33)$$

If

$$c_{21} < \frac{1}{2}c_{12} + \frac{1}{2}c_{11}$$

the value of private information for player 2 is negative. On the other hand, in this situation, the gain for player 1 is, according to (4-13) and (4-23)

$$J_1^P - J_1^0 = \frac{1}{2}(c_{12} - c_{22}) > 0 \quad (4-34)$$

Thus, for the class of symmetric games considered in this paper, we have a case where the value of private information is negative for the informed player and the uninformed player benefits from the situation! This phenomenon never occurs for this class of games in the classical Bayesian framework, where if the value of private information is negative for the informed player the uninformed player cannot benefit either [15].

Even more surprising in this case is the fact that the informed player wants to use his private information, whereas the uninformed player wishes that the informed player acted as if he were not informed!!

The reason for all these counterintuitive results and the differences between the subjective game and the classical Bayesian game is that each player evaluates the game as well as the behavior of his opponent in the game in terms of his own model and acts accordingly. Such subjective evaluations lead to behavior which would never occur in the classical Bayesian formulation as evidenced by the previous analysis.

One issue that naturally arises in these games is the following: How do the players involved in the game interpret its outcome? Do they realize that they have different models? If neither player is informed about the outcome of the chance move, then player 1 plays the game described by (4-20) and player 2 plays the game described by (4-17). In this situation, player 1 expects that player 2 will use strategy σ and player 2 expects that player 1 will use strategy μ . At the end of the game, each player finds out that the outcome is the opposite of what he expected. Since each player assumes that his opponent is rational, at the end of the game both players conclude that they have different models. Similar phenomena occur if one of the players is either secretly or privately informed.

In the case of secret information, the secretly informed player discovers at the end of the game that his opponent's perception of the game is different from his. On the other hand, the uninformed player may 1) never discover that his opponent has a different perception of the game or, 2) not be able to interpret his opponent's move in terms of his own model in which case he can conclude that 1) either his opponent has a different model of the game or (most likely), (ii) his opponent has secret information.

In the case of private information, the uninformed player is not in a position to discover at the end of the game that his opponent has a different view of the game. The informed player may or may not (depending on whether (4-12) holds) discover at the end of the game that he and his opponent have inconsistent beliefs about the game.

Note that if both $p, q > 1/2$ or $p, q < 1/2$, the players never discover the differences in their models.

4.4 Summary

In this section, we presented and analyzed a simple class of two person non-cooperative, non-zero sum one-stage subjective games of incomplete information. We showed how the inconsistent beliefs of the players lead to results which are counterintuitive and different from the results of the classical Bayesian game, and how private or secret information is differently evaluated by each player.

One important issue that has not been discussed so far is the following: Are the differences between the two players amplified or smoothed out if the game is repeated over and over? This issue will not be discussed here due to space limitation. The analysis of the repeated subjective game of incomplete information appears in [16].

5. SUMMARY - CONCLUSIONS

In this paper, we presented a formulation of subjective games. This formulation relaxes two of the key assumptions upon which game theory was originally developed. 1) the assumption that all players have the same model of the game and 2) the assumption that all the rules of the game are common knowledge to the players. To illustrate the differences between our formulation and previous formulations of games, we restricted attention to a simple specific class of games of incomplete information. We showed that even for this simple class of games, the inconsistent beliefs of the players lead to results which are counterintuitive and different from the results of the corresponding classical Bayesian games.

REFERENCES

1. Aumann, R.J. and M. Maschler, "Some Thoughts on the Minimax Principle." Management Science, Vol. 18, No. 5, January 1972.
2. Ho, Y.C., "On the Minimax Principle and Zero Sum Stochastic Differential Games." Technical Report No. 630, Harvard University, Cambridge, MA, April 1972.
3. Kreps, D.M. and R. Wilson, "Sequential Equilibria," Stanford University Report, December 1980.
4. Wilson, R., "Perfect Equilibria and Sequential Rationality." presented at Luminy Conference on Games with Incomplete Information, Sept. 1981.
5. Selten, R., "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games." International Journal of Game Theory, Vol. 4, No. 1, pp. 25-55.
6. Selten, R., "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit." Zeitschrift für die gesamte Staatswissenschaft, Vol. 12, 301-324.
7. Harsanyi, J.C., "Games of Incomplete Information Played by Bayesian Players," Parts I-III, Management Science, Vol. 14, No. 3, No. 5, No. 7, 1967-1968.
8. Aumann, R.J., M. Maschler et al. "Game Theoretic Aspects of Gradual Disarmament." In: Mathematica 1966-1968
9. Sorin, S. "An Introduction to Two-Person Zero-Sum Repeated Games of Incomplete Information," Technical Report No. 312, Institute for Mathematical Studies in the Social Sciences, Stanford University, May 1980.
10. Hart, S., "Non-Zero Sum Two-Person Repeated Games with Incomplete Information," Technical Report No. 367, Institute for Mathematical Studies in the Social Sciences, Stanford University, February 1982.
11. Aumann, R.J., "Acceptable Points in General Cooperative n-Person Games," Contributions to the Theory of Games, Vol. IV, Annals of Mathematical Studies, No. 40, Princeton: Princeton University Press, pp. 287-324
12. VonNeumann, J. and Morgenstern, O., Theory of Games and Economic Behavior, J. Wiley, 1944.
13. Savage, L., The Foundations of Statistics, Dover 1972.
14. Teneketzis, D. and P. Varaiya, "Asymptotic Agreement in Distributed Estimation with Inconsistent Beliefs", ALPHATECH Technical Report TR-172, September 1983.
15. Levine, P., and J.P. Ponsard, "The Values of Information in Some Non-zero Sum Games." International Journal of Game Theory, Vol. 6, No. 4, pp. 221-229.
16. Teneketzis, D., and D. A. Castanon, "Informational Aspects of a Class of Subjective Games of Incomplete Information," ALPHATECH Technical Report TR-189, November 1983.

Dup^e

ON THE COMPLEXITY OF DISTRIBUTED DECISION PROBLEMS*

by

John Tsitsiklis and Michael Athans
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

ABSTRACT

We study the computational complexity of finite versions of the simplest and fundamental problems of distributed decision making and we show that, apart from a few exceptions, such problems are hard (NP-complete, or worse). Some of the problems studied are the well-known team decision problem, the distributed hypothesis testing problem, as well as the problem of designing a communications protocol that guarantees the attainment of a prespecified goal with as little communications as possible. These results indicate the inherent difficulty of distributed decision making, even for very simple problems, with trivial centralized counterparts and suggest that optimality may be an elusive goal of distributed systems.

1. Introduction and Motivation

In this paper we formulate and study certain simple decentralized problems. Our goal is to formulate problems which reflect the inherent difficulties of decentralization; that is, any difficulty in this class of problems is distinct from the difficulty of corresponding centralized problems. This is accomplished by formulating decentralized problems whose centralized counterparts are either trivial or vacuous.

One of our goals is to determine a boundary between "easy" and "hard" decentralized problems. Our results will indicate that the set of "easy" problems is relatively small.

All problems to be studied are imbedded in a discrete framework; the criteria we use for deciding whether a problem is difficult or not come from complexity theory [Garey and Johnson, 1979; Papadimitriou and Steiglitz, 1982]: following the tradition of complexity theory, problems that may be solved by a polynomial algorithm are considered easy; NP-complete, or worse, problems are considered hard.* However, an NP-completeness result does not close a subject, but is rather as a result which can guide research: further research should focus on special cases of the problem or on approximate versions of the original problem.

The main issue of interest in decentralized systems may be loosely phrased as "who should communicate to whom, what, how often etc." From a purely logical point of view, the first question that has to be raised is "are there any communication necessary?" Any further questions deserve to be studied only if we come to the conclusion that communications are indeed necessary.

The subject of Section 2 is to characterize the inherent difficulty of the problem of deciding whether any communications are necessary, for a given situation. We adopt the following approach: a decentralized system exists in order to accomplish a certain goal which is externally specified and well-known. A

set of processors obtain (possibly conflicting) observations on the state of the environment. Each processor has to make a decision, based on his own observation. However, for each state of the environment, only certain decisions accomplish the desired goal. The question "are there any communications necessary?" may be then reformulated as "can the goal be accomplished, with certainty, without any communications?" We show that this problem is, in general, a hard one.

We then impose some more structure on the problem, by assuming that the observations of different processors are related in a particular way. The main issue that we address is "how much structure is required so that the problem is an easy one?" and we try to determine the boundary between easy and hard problems.

In Section 3 we formulate a few problems which are related to the basic problem of Section 2 and discuss their complexity.

In Section 4 we study a particular (more structured) decentralized problem - the problem of decentralized hypothesis testing - on which there has been some interest recently, and characterize its difficulty.

Suppose that it has been found that communications are necessary. The next question of interest is "what is the least amount of communications needed?" This problem (Section 5) is essentially the problem of designing an optimal communications protocol; it is again a hard one and we discuss some related issues.

In Section 6 we present our conclusions and discuss the conceptual significance of our results. These conclusions may be summarized by saying that:

- a) Even the simplest (exact) problems of decentralized decision making are hard.
- b) Allowing some redundancy in communications, may greatly facilitate the (off-line) problem of designing a decentralized system.
- c) Practical communications protocols should not be expected to be optimal, as far as minimization of the amount of communications is concerned.

Some of the results of this paper appear in [Papadimitriou and Tsitsiklis, 1983] and (almost) all proofs may be found in [Tsitsiklis, 1983].

2. A Problem of Silent Coordination

In this section we formulate and study the problem whether a set of processors with different information may accomplish a given goal -with certainty- without any communications.

Let $\{1, \dots, M\}$ be a set of processors. Each processor, say processor i , obtains an observation y_i which comes from a finite set Y_i of possible observations. Then, processor i makes a decision u_i which belongs to a finite set U_i of possible decisions, according to a rule.

* Research supported by ONR under contract ONR/N00014-77-C-0532 (NR-041-519).

* One way of viewing NP-complete problems, is to say that they are effectively equivalent to the Traveling Salesman problem, which is well-known to be algorithmically hard.

$$u_i = \partial_i(y_i), \quad (2.1)$$

where ∂_i is some function from Y_i into U_i . The M -tuple (y_1, \dots, y_M) is the total information available; so it may be viewed as the "state of the environment." For each state of the environment, we assume that only certain M -tuples (u_1, \dots, u_M) of decision accomplish a given, externally specified, goal. More precisely, for each $(y_1, \dots, y_M) \in Y_1 \times \dots \times Y_M$ we are given a set $S(y_1, \dots, y_M) \subset U_1 \times \dots \times U_M$ of satisficing decisions.

(So, S may be viewed as a function from

$$Y_1 \times Y_2 \times \dots \times Y_M \text{ into } 2^{U_1 \times \dots \times U_M}.$$

The problem to be studied, which we call "distributed satisficing problem" (after the term introduced by H. Simon [1980]) may be described formally as follows:

Distributed Satisficing (DS): Given finite sets Y_1, \dots, Y_M , U_1, \dots, U_M and a function $S: Y_1 \times \dots \times Y_M \rightarrow$

$2^{U_1 \times \dots \times U_M}$, are there functions $\partial_i: Y_i \rightarrow U_i$, $i=1, 2, \dots, M$, such that

$$(\partial_1(y_1), \dots, \partial_M(y_M)) \in S(y_1, \dots, y_M), \quad \forall (y_1, \dots, y_M) \in Y_1 \times \dots \times Y_M \quad (2.2)$$

Remarks:

1. We are assuming that the function S is "easily computable," for example, it may be given in the form of a table.
2. The centralized counterpart of DS would be to allow the decision u_i of each agent depend on the entire set (y_1, \dots, y_M) of observations; so, ∂_i would be a function from $Y_1 \times \dots \times Y_M$ into U_i . (This corresponds to a situation in which all processors share the same information.). Clearly, then, there exist satisfactory (satisficing) functions $\partial_i: Y_1 \times \dots \times Y_M \rightarrow U_i$, if and only if $S(y_1, \dots, y_M) \neq \emptyset$, $\forall (y_1, \dots, y_M) \in Y_1 \times \dots \times Y_M$. Since S is an "easily computable" set as a function of its arguments, we can see that the centralized counterpart of DS is a trivial problem. So, any difficulty inherent in DS is only caused by the fact that information is decentralized.
3. A "solution" for the problem DS cannot be a closed-form formula which gives an answer 0 (no) or 1 (yes). Rather, it has to be an algorithm, a sequence of instructions, which starts with the data of the problem $(Y_1, \dots, Y_M, U_1, \dots, U_M, S)$ and eventually provides the correct answer. Accordingly, the difficulty of the problem DS may be characterized by determining the place held by DS in the complexity hierarchy. For definitions related to computational complexity and the methods typically used, the reader is referred to [Garey and Johnson, 1979; Papadimitriou and Steiglitz, 1982].
4. If, for some i , the set U_i is a singleton, processor i has no choice, regarding his decision and, consequently, the problem is equivalent to a problem in which processor i is absent. Hence, without loss of generality, we only need to study instances of DS in which $|U_i| \geq 2$, $\forall i$.
5. We believe that the problem DS captures the essence of coordinated decision making with decentralized information and without communications (silent coordination).

Some initial results on DS are given by the

following:

Theorem 2.1:

- a) The problem DS with two processors ($M=2$) and restricted to instances for which the cardinality of the decision sets is 2 ($|U_i|=2$, $i=1,2$) may be solved in polynomial time.
- b) The problem DS with two processors ($M=2$) is NP-complete, even if we restrict to instances for which $|U_1|=2$, $|U_2|=3$.
- c) The problem DS with three (or more) processors ($M \geq 3$) is NP-complete, even if we restrict to instances for which $|U_i|=2$, $\forall i$.

Theorem 2.1 states that the problem DS is, in general, a hard combinatorial problem, except for the special case in which there are only two processors and each one has to make a binary decision. It should be noted that the difficulty is not caused by an attempt to optimize with respect to a cost function, because no cost function has been introduced. In game theoretic language, we are faced with a "game of kind" rather than a "game of degree."

We will now consider some special cases (which reflect the structure of typical practical problems) and examine their computational complexity, trying to determine the dividing line between easy and hard problems. From now on we restrict our attention to the case in which there are only two processors. Clearly, if a problem with two processors is hard, the corresponding problem with three or more processors cannot be easier.

We have formulated above the problem DS so that all pairs $(y_1, y_2) \in Y_1 \times Y_2$ are likely to occur. So, the information of different processors is completely unrelated; their coupling is caused only by the structure of the satisficing sets $S(y_1, y_2)$. In most practical situations, however, information is not completely unstructured: when processor 1 observes y_1 , he is often able to make certain inferences about the value of the observation y_2 of the other processor and exclude certain values. We now formalize these ideas:

Definition: An Information Structure I is a subset of $Y_1 \times Y_2$. We say that an information structure I has degree (D_1, D_2) (D_1, D_2 are positive integers) if

- (i) For each $y_1 \in Y_1$ there exist at most D_1 distinct elements of Y_2 such that $(y_1, y_2) \in I$.
 - (ii) For each $y_2 \in Y_2$ there exist at most D_2 distinct elements of Y_1 such that $(y_1, y_2) \in I$.
 - (iii) D_1, D_2 are the smallest integers satisfying (i), (ii).
- An information structure I is called classical if $D_1=D_2=1$; nested if $D_1=1$ or $D_2=1$.

We now interpret this definition: The information structure I is the set of pairs (y_1, y_2) of observations that may occur together. If I has degree (D_1, D_2) processor 1 may use his own observation to decide which elements of Y_2 may have been observed by processor 2.

In particular, he may exclude all elements except for D_1 of them. The situation faced by processor 2 is symmetrical.

If $D_1=1$ and processor 1 observes y_1 , there is only one possible value for y_2 . So, processor 1 knows the observation of processor 2. (The converse is true when $D_2=1$). This is called a nested information structure because the information of one processor contains the information of the other.

When $D_1=D_2=1$, each processor knows the observation of the other; so, their information is essentially

shared.

Since pair (y_1, y_2) not in I cannot occur, there is no meaning in requiring the processors to make compatible decisions if (y_1, y_2) were to be observed. This leads to the following version of the problem DS:

DSI: Given finite sets Y_1, Y_2, U_1, U_2 , $I \subset Y_1 \times Y_2$ and a function $S: I \rightarrow 2^{U_1 \times U_2}$, are there functions $\delta_i: Y_i \rightarrow U_i$, $i=1,2$, such that

$$(\delta_1(y_1), \delta_2(y_2)) \in S(y_1, y_2), \forall (y_1, y_2) \in I? \quad (2.3)$$

Note that any instance of DSI is equivalent to an instance of DS in which $S(y_1, y_2) = U_1 \times U_2$, $\forall (y_1, y_2) \in I$.

That is, no compatibility restrictions are placed on the decisions of the two processors, for those (y_1, y_2) that cannot occur.

We now proceed to the main result of this Section:

Theorem 3.2.2:

a) The problem DSI restricted to instances satisfying any of the following:

- (i) One or more of $|U_1|, |U_2|, D_1, D_2$ is equal to 1.
- (ii) $|U_1| = |U_2| = 2$,
- (iii) $D_1 = D_2 = 2$,
- (iv) $D_1 = |U_2| = 2$, (or $D_2 = |U_1| = 2$)

may be solved in polynomial time.

b) The problem DSI is NP-complete even if we restrict to instances for which

$$|U_1| = D_1 = 3, |U_2| = D_2 = 2$$

The result concerning the case $D_1=1$ or $D_2=1$ is not surprising. It is well-known that nested information structures may be exploited to solve otherwise difficult decentralized problems. But except for the case $D_1=D_2=2$ (which is sort of a boundary) the absence of nestedness makes decentralized problems computationally hard. Our result gives a precise meaning to the statement that non-nested information structures are much more difficult to handle than nested ones.

Theorem 3.2.2 shows that even if D_1, D_2 are held constant, the problem DSI is, in general, NP-complete. There is, however, a special case of DSI, with D_1, D_2 constant, for which an efficient algorithm of the dynamic programming type is possible:

Theorem 3.2.3: Let $Y_1 = \{1, \dots, m\}$, $Y_2 = \{1, \dots, n\}$ and suppose that $|i-j| \leq D$, $\forall (i,j) \in I$. Then, if D is held constant, DSI may be solved in polynomial time.

Remark: In fact, the conclusion of Theorem 3.2.3 remains true if we assume $m=n$ and we replace the condition $|i-j| \leq D$ by the weaker condition $|i-j| \pmod n \leq D$. The proof consists of a small modification of the preceding one.

The condition $|i-j| \leq D$, $\forall (i,j) \in I$ is fairly natural in certain applications. For example, suppose that the observations y_1 and y_2 are noisy measurements of an unknown variable x ($y_i = x + w_i$) where the noises w_i are bounded: $|w_i| \leq D/2$.

The condition $|i-j| \pmod n \leq D$ may also arise if the observations y_1, y_2 are noisy measurements of some unknown angle: $y_i = \theta + w_i$

3. Related Problems

In this Section we define and discuss briefly a few more combinatorial problems relevant to decentralized decision making. All of them will be seen to be harder than problem DS of the last section (i.e. they contain DS as a special case) and are, therefore NP-hard (that is, NP-complete, or worse).

The best known static decentralized problem is the team decision problem [Marschak and Badner, 1972] which admits an elegant solution under linear quadratic assumptions. Its discrete version is the following:

TDP (Team Decision Problem): Given finite sets Y_1, Y_2, U_1, U_2 , a probability mass function $p: Y_1 \times Y_2 \rightarrow Q$, and a cost function $c: Y_1 \times Y_2 \times U_1 \times U_2 \rightarrow N$, find decision rules $\delta_i: Y_i \rightarrow U_i$, $i=1,2$ which minimize the expected cost

$$J(\delta_1, \delta_2) = \sum_{y_1 \in Y_1} \sum_{y_2 \in Y_2} c(y_1, y_2, \delta_1(y_1), \delta_2(y_2)) p(y_1, y_2)$$

Let $S(y_1, y_2) = \{(u_1, u_2) \in U_1 \times U_2 : c(y_1, y_2, u_1, u_2) = 0\}$. If we solve TDP, we have effectively answered the question whether there exist δ_1, δ_2 such that $J(\delta_1, \delta_2) = 0$. This is equivalent to the question whether there exist satisfying decision rules (with the satisfying sets $S(y_1, y_2)$ defined as above). Therefore, TDP is harder than DS:

Proposition 3.1: The discrete team decision problem is NP-hard, even if the range of the cost function c is $\{0,1\}$.

Instead of trying to "satisfice" for every pair of observations $(y_1, y_2) \in Y_1 \times Y_2$, it may be more appropriate to impose a probability mass function $\phi: Y_1 \times Y_2 \rightarrow Q$ and try to maximize the probability of satisficing. This leads to the next problem:

MPS (Maximize Probability of Satisficing): Given finite sets Y_1, Y_2, U_1, U_2 , a probability mass function

$p: Y_1 \times Y_2 \rightarrow Q$ and a function $S: Y_1 \times Y_2 \rightarrow 2^{U_1 \times U_2}$, find decision rules $\delta_i: Y_i \rightarrow U_i$, $i=1,2$, which maximize the probability of satisficing $J(\delta_1, \delta_2) = \Pr((\delta_1(y_1), \delta_2(y_2)) \in S(y_1, y_2))$.

We now take a slightly different point of view. Suppose that communications are allowed, so that the processors may always make satisficing decisions by communicating (assuming that $S(y_1, y_2) \neq \emptyset$, $\forall (y_1, y_2) \in Y_1 \times Y_2$). Suppose, however, that communications are very expensive, so that we are interested in a scheme which guarantees satisficing with a minimum amount of communications. We will assume that if one of the processors initiates a communication, all their information will be exchanged at unity cost. (For a more refined way of counting the amount of communications, see Section 3.5.)

MPC (Minimize Probability of Communications): Given finite sets Y_1, Y_2, U_1, U_2 a probability mass function

$p: Y_1 \times Y_2 \rightarrow Q$ and a function $S: Y_1 \times Y_2 \rightarrow 2^{U_1 \times U_2}$, find decision rules $j_i: Y_i \rightarrow U_i \cup \{c\}$, $i=1,2$, which minimize the probability $\Pr(\delta_1(y_1) = c \text{ or } \delta_2(y_2) = c)$ of communicating subject to the constraint

If $\{\partial_1(y_1) \neq C \text{ and } \partial_2(y_2) \neq C\}$ then $(\partial_1(y_1), \partial_2(y_2)) \in S(y_1, y_2)$.

The proof of the following is trivial:

Proposition 3.2: The problems MPS and APC are NP-hard.

In fact, we also have:

Proposition 3.3: The problems TDP (with a zero-one cost function) and MPS are NP-hard, even if $|U_1| = |U_2| = 2$.

We could also define dynamic versions of DS or of the team problem, in a straightforward way [Tenney, 1983]. Since dynamic problems cannot be easier than static ones, they are automatically NP-hard.

4. Decentralized Hypothesis Testing

A basic problem in decentralized signal processing, which has attracted a fair amount of attention recently, is the problem of decentralized hypothesis testing [Tenney and Sandell, 1981; Ekchian, 1982; Ekchian and Tenney, 1982; Kushner and Pacut, 1982; Lauer and Sandell, 1983]. A simple version of the problem, involving only two processors and two hypotheses may be described as follows:

Two processors S_1 and S_2 receive observations $y_1 \in Y_1$, $y_2 \in Y_2$, respectively, where Y_i is the set of all possible observations of processor i . (Figure 1). There are two hypotheses H_0 and H_1 on the state of the environment, with prior probabilities p_0 and p_1 respectively. For each hypothesis H_i , we are also given the joint probability distribution $P(y_1, y_2 | H_i)$ of the observations, conditioned on the event that H_i is true.

Upon receipt of y_i , processor S_i evaluates a message $u_i \in \{0, 1\}$ according to the rule $u_i = \partial_i(y_i)$, where $\partial_i: Y_i \rightarrow \{0, 1\}$. Then, u_1 and u_2 are transmitted to a




Figure 1: A Scheme for Decentralized Hypothesis Testing.

central processor (fusion center) which evaluates $u_0 = u_1 \cap u_2$ and declares hypothesis H_0 to be true if $u_0 = 0$, H_1 if $u_0 = 1$. (So, we essentially have a voting scheme). The problem is to select the functions ∂_1 , ∂_2 so as to minimize the probability of accepting the wrong hypothesis. (More general performance criteria may be also considered).

Most available results assume that

$$P(y_1, y_2 | H_i) = P(y_1 | H_i)P(y_2 | H_i), \quad i=1,2, \quad (4.1)$$

which states that the observations of the two processors are independent, when conditioned on either hypothesis.* In particular, it has been shown [Tenney and Sandell, 1981] that the optimal decision rules ∂_i are given in terms of thresholds for the likelihood

ratios $\frac{p_0 P(H_0 | y_1)}{p_1 P(H_1 | y_1)}$. The optimal thresholds for

the two sensors are coupled through a system of equations which give necessary conditions of optimality. (These equations are precisely the person-by-person optimality conditions). Few analytical results are available when the conditional independence assumption is removed [Lauer and Sandell, 1983]. The approach of this section is aimed at explaining this status of affairs, by focusing on discrete (and finite) versions of the problem.

We first have:

Theorem 4.1: If Y_1, Y_2 are finite sets and (4.1) holds, then optimal choices for ∂_1, ∂_2 may be found in polynomial time.

So, under the conditional independence assumption, decentralized hypothesis testing is a computationally easy problem. Unfortunately, this is not the case when the independence assumption is relaxed. Our main result (Theorem 4.2) states that (with Y_1, Y_2 finite

sets), decentralized hypothesis testing is a hard combinatorial problem (NP-hard). This is true even if we restrict to the special case where perfect detection (zero probability of error) is possible for the corresponding centralized hypothesis testing problem. Although this is in some sense a negative result, it is useful because it indicates the direction in which future research on this subject should proceed:

Instead of trying to find efficient exact algorithms, research should focus on approximate algorithms, or exact algorithms for problems with more structure than that assumed here. Moreover, our result implies that any necessary conditions for optimality to be developed are likely to be deficient in one of two respects:

- a) Either there will be a very large number of decision rules satisfying these conditions.
- b) Or, it will be hard to find decision rules satisfying these conditions.

In particular, optimal decision rules are not given in terms of thresholds on likelihood ratios.

Of course, there remains the question whether efficient approximate algorithms exist for the general decentralized hypothesis testing problem, or whether we must again restrict to special cases of the problem.

We now present formally the problem to be analyzed.

DHT: (Decentralized Hypothesis Testing, Restricted to Instances for which Perfect Centralized Detection is Possible).

We are given finite sets Y_1, Y_2 ; a rational number k ; a rational probability mass function $p: Y_1 \times Y_2 \rightarrow \mathbb{Q} \cap [0, 1]$; a partition

* Such an assumption is reasonable in problems of detection of a known signal in independent noise, but is typically violated in problems of detection of an unknown signal.

$\{A_0, A_1\}$ of $y_1 \times y_2$.^{*} Do there exist $\delta_1, \delta_2 \in \{0,1\}$, $\delta_2, \delta_2 \in \{0,1\}$ such that $J(\delta_1, \delta_2) \leq k$, where

$$J(\delta_1, \delta_2) = \sum_{(y_1, y_2) \in A_0} P(y_1, y_2) \delta_1(y_1) \delta_2(y_2) + \sum_{(y_1, y_2) \in A_1} P(y_1, y_2) [1 - \delta_1(y_1) \delta_2(y_2)]? \quad (4.2)$$

Remarks: 1. If we let $k=0$, then DHT is a special case of problem DS (Section 2), with $|U_1|=|U_2|=2$, and is polynomially solvable, according to Theorem 3.2.1. In general DHT is a special case of MTS and TDF (Section 3.3) with $|U_1|=|U_2|=2$. Consequently, Theorem 4.2

below proves Proposition 3.3.

2) Clearly, the optimization problem (Minimize $J(\delta_1, \delta_2)$ with respect to δ_1, δ_2) cannot be easier than DHT.

Since DHT will be shown to be NP-complete, it follows that the above optimization problem is NP-hard.

3) In DHT, as defined above, we are only considering instances for which perfect centralized detection is possible: Think of H_0 as being the hypothesis that

$(y_1, y_2) \in A_0$, and H_1 as being the hypothesis that $(y_1, y_2) \in A_1$. Certainly, if a processor knows both y_1, y_2 , the true hypothesis may be found with certainty.

For the decentralized problem, the cost function $J(\delta_1, \delta_2)$ is easily seen to be the probability of error.

4) The result to be obtained below remains valid if the fusion center uses different rules for combining the messages it receives (e.g. $u_0 = (u_1 \vee u_2)$), or if we leave the combining rule unspecified and try to find an optimal combining rule.

Theorem 4.2: DHT is NP-complete.

5. On Designing Communications Protocols

Suppose that we are given an instance of the distributed satisficing problem (DS) and that it was concluded that unless the processors communicate, satisficing cannot be guaranteed for all possible observations. Assuming that communications are allowed (but are costly), we have to consider the problem of designing a communications protocol: what should each processor communicate to the other, and at what order? Moreover, since communications are costly, we are interested in a protocol which minimizes the total number of binary messages (bits) that have to be communicated. (The word "bits" above does not have the information theoretic meaning.)

Before proceeding, we must make more precise the notion of a communication protocol and of the number of bits than guarantee satisficing.

Given an instance $D = (Y_1, Y_2, U_1, U_2, I, S)$ of the problem DS1 we will say that:

There is a protocol which guarantees satisficing with 0 bits of communications, if D is a YES instance of the problem DS1. (That is, if there exist satisficing decision rules, involving no communications.)

We then proceed inductively:

There is a protocol which guarantees satisficing with K bits of communications (KE N), if for some $i \in \{1, 2\}$ (say, $i=1$) there is a function $m: Y_1 \rightarrow \{0,1\}$, such that for each of the instances $D' = (Y_1 \cap m^{-1}(0), Y_2, U_1, U_2, I \cap (Y_1 \cap m^{-1}(0)) \times Y_2, S)$ and

^{*} That is $A_0 \cap A_1 = Y_1 \times Y_2$ and $A_0 \cap A_1 = \emptyset$.

$D'' = (Y_1 \cap m^{-1}(1), Y_2, U_1, U_2, I \cap (Y_1 \cap m^{-1}(1)) \times Y_2, S)$ there is a protocol which guarantees satisficing with not more than $K-1$ bits of communications. (Here $m^{-1}(i) = \{y_1 \in Y_1 : m(y_1) = i\}$.)

The envisaged sequence of events behind this definition is the following: Each processor observes his measurement $y_i \in Y_i$, $i=1,2$. Then, one of the processors, say processor 1, transmits a message $m(y_1)$, with a single bit to the other processor. From that point on, it has become common knowledge that $y_1 \in Y_1 \cap m^{-1}(y_1)$;

therefore, the remaining elements of Y_1 may be ignored.

We can now state formally the problem of interest:

MBS (Minimum bits to suffice): Given an instance D of DS1 and KE N , is there a protocol which guarantees satisficing with not more than K bits of communications?

By definition, MBS with $K=0$ is identical to the problem DS1. Moreover, MBS with K arbitrary cannot be easier than MBS with $K=0$ (which is a special case). Therefore, MBS is, in general NP-hard. Differently said, problems involving communications are at least as hard as problems involving no communications.

We have seen in Section 2 that when $|U_1|=|U_2|=2$,

DS1 may be solved in polynomial time. Therefore, MBS with $K=0$, $|U_1|=2$, $|U_2|=2$ is polynomially solvable.

However, for arbitrary K , this is no longer true:

Theorem 5.1: MBS is NP-complete, even if $|U_1|=|U_2|=2$ and even if we restrict to instances for which, for any $(y_1, y_2) \in I$, either $S(y_1, y_2) = \{(0,0)\}$ or $S(y_1, y_2) = \{(1,1)\}$.

The above theorem proves a conjecture of A. Yao [Yao, 1979]. The proof was mainly constructed by C. Papadimitriou and may be found in [Papadimitriou and Tsitsiklis, 1982].

We should point out that the special case referred to in Theorem 5.1 concerns the problem of distributed function evaluation: we are given a Boolean function $f: Y_1 \times Y_2 \rightarrow \{0,1\}$ and we require that both agents (processors) eventually determine the value of the function (given the observation -input (y_1, y_2)), by exchanging a minimum number of bits. In our formalism, $S(y_1, y_2) = \{(0,0)\}$ if $f(y_1, y_2) = 0$ and $S(y_1, y_2) = \{(1,1)\}$ if $f(y_1, y_2) = 1$.

In Section 2 we had investigated the complexity of DS1 by restricting to instances for which the set I had constant degree (D_1, D_2) . This may be done, in principle, for MBS, as well, but no results are available, except for the simple case in which $D_1 = D_2 = 2$.

In fact, when $D_1 = D_2 = 2$ each processor may transmit his information to the other agent by communicating a single binary message and, for this reason, we have:

Proposition 5.1: MBS restricted to instances for which $D_1 = D_2 = 2$ may be solved in polynomial time. Moreover, an optimal protocol requires transmission of at most two binary messages, one from each processor.

When (D_1, D_2) is larger than $(2,2)$, there is not much we can say about optimal protocols. However, it is easy to verify that there exist fairly simple non-optimal protocols (which may be calculated in polynomial time) which involve relatively small amounts of communication. This is because:

Proposition 5.2: Suppose that I has degree (D_1, D_2) and that $S(y_1, y_2) \neq \emptyset$, $\forall (y_1, y_2) \in I$. Then information may be centralized (and therefore satisficing is guaranteed) by means of a protocol requiring communication of at

most $\lceil \log_2(D_1 D_2) \rceil$ binary messages by each processor. Moreover, such a protocol may be constructed in time $O((|Y_1| + |Y_2|)(|Y_1| + |Y_2|))$. (Here $\lceil x \rceil$, $x \in \mathbb{R}$, stands for the smallest integer larger than x .)

Remark 1: It might be tempting to guess that processor 1 (respectively 2) needs to communicate only $\lceil \log_2 D_2 \rceil$ (respectively $\lceil \log_2 D_1 \rceil$) bits, but this is not true, as can be seen from fairly simple examples.

6. Conclusions

We summarize here the main conclusions of this paper.

Even if a set of processors have complete knowledge of the structure of a distributed decision making problem and the desired goal; even if the corresponding centralized problem is trivial; even if all relevant sets are finite, a satisficing decision rule that involves no on-line communications may be very hard to find, the corresponding problem being, in general, NP-complete. There are many objections to the idea that NP-completeness is an unequivocal measure of the difficulty of a problem, because it is based on a worst case analysis, whereas the average performance of an algorithm might be a more adequate measure; moreover NP-hard optimization problems may have very simple approximate algorithms. However, NP-complete problems are often characterized by the property that any known algorithm is very close to systematic exhaustive search; they do not possess any structure to be exploited.

Concerning the problem DS, and its variations, we may reach the following specific conclusions: No simple algorithm could solve DS. Given that communications would be certainly required for those instances of DS that possess no satisficing decision rules, it would not be a great loss if we allowed the processors to communicate even for some instances of DS for which this would not be necessary. Even if these extra communications -being redundant- do not lead to better decisions, they may greatly facilitate the decision process and -from a practical point of view - remove some load from the computing machines employed.

Concerning the problem of distributed hypothesis testing, we have shown that it becomes hard, once a simplifying assumption of conditional independence is removed. This explains why no substantial progress on this problem had followed the work of Tenney and Sandell [1982].

From a more general perspective, we are in a position to say that the basic (and the simplest) problems of decentralized decision making are hard, in a precise mathematical sense. Moreover, their difficulty does not only arise when one is interested in optimality. Difficulties persist even if optimality is replaced by satisficing. As a consequence, further research should focus on special cases and easily solvable problem as well as on approximate versions of the original problems.

In cases where communications are necessary (but costly) there arises naturally the problem of designing a protocol of communications. Unfortunately, if this problem is approached with the intention to minimize the amount of communications that will guarantee the accomplishment of a given goal, we are again led to intractable combinatorial problems. Therefore, practical communications protocols can only be designed on a "good" heuristic or ad-hoc basis, and they should not be expected to be optimal; approximate optimality is probably a more meaningful goal. Again, allowing some redundancy in on-line communications may lead to substantial savings in off-line computations.

6. References

- Ekchian, L.K., (1982) "Distributed Detection and Communications Problems," Ph.D. Thesis, Dept. of EECS, MIT, Cambridge, MA.
- Ekchian, L.K., Tenney, R.R., (1982), "Detection Networks," Technical Report LIDS-P-1191, Laboratory for Information and Decision Systems, MIT, Cambridge, MA.
- Garey, M.R., D.S. Johnson, (1979), Computers and Intractability: A Guide to the Theory of NP-Completeness, W.H. Freeman and Co., San Francisco.
- Kushner, H.J., A. Pacut, (1982), "A Simulation Study of a Decentralized Detection Problem," IEEE Trans. on Auto. Control, Vol. AC-27, No. 5, pp. 1116-1119.
- Lauer, G.S., N.R. Sandell, Jr. (1983), "Distributed Detection of Signal Waveforms in Additive Gaussian Observation Noise," Technical Paper 160, Alphatech Inc., Burlington, Mass.
- Marschak, J., R. Radner, (1972), The Economic Theory of Teams, Yale University Press, New Haven, 1972.
- Papadimitriou, C.H., K. Steiglitz, (1982), Combinatorial Optimization: Algorithms and Complexity, Prentice-Hall.
- Papadimitriou, C.H., J. Tsitsiklis, (1983), "On the Complexity of Designing Distributed Protocols." to appear in Information and Control.
- Simor, H.A., (1980), The Sciences of the Artificial, MIT Press, Cambridge, MA.
- Tenney, R.R., (1983), "Optimal Decentralized Control of Finite Nondeterministic Systems," Proceedings of the 1983 American Control Conference.
- Tenney, R.R., N.R. Sandell, Jr., (1981), "Detection with Distributed Sensors," IEEE Trans. on Aerospace and Electronic Systems, Vol. AES-17, No. 4, pp. 501-509.
- Tsitsiklis, J.N., (1982), "Problems in Distributed Decision Making," Technical Report, Lab. for Information and Decision Systems, MIT, in preparation.
- Yao, A.C.-C., (1979), "Some Complexity Questions Related to Distributed Computing," Proceedings of the 11th Annual Symposium on Theory of Computing, pp. 209-213.

Robert R. Tenney²

Massachusetts Institute of Technology, Cambridge, Mass. 02139

Summary

A control problem with a system modeled as a nondeterministic finite state machine is considered. Several agents seek to optimize the behavior of the system under a minimax criterion, with each agent having different information about the system state. The nondeterministic model of uncertainty, combined with the minimax criterion, lead to equivalence relations on the past input/output histories of each agent which generate simple sufficient statistics for the optimal control laws. This sheds light on the basic nature of decentralized control and permits complete solution of a particular class of problems.

I. Introduction

A system to be controlled is assumed to be described as a finite state, nondeterministic automaton with inputs supplied by several control agents. Each agent receives an observation at each discrete time step which indicates a set in which the current state must lie. Each observation, coupled with knowledge of the system structure, can lead to inferences about the past system behavior, hence about the past observations of other agents and thus predictions of other agents' decisions. This simultaneous interweaving of inference by the agents, as each deduces the potential actions of others, and the deductions of others about itself, etc. leads to some of the complexity of the analytical process for general decentralized problems.

Based on this model structure, an approach to addressing the problem of designing the optimal (in the appropriate worst-case sense) decision rules has been developed. The approach is based on the identification of the set of sufficient statistics for each agent to use and the dynamic relations between them; these sufficient statistics are no more than the intertwined deductions of the agents about each other truncated at the point where they are no longer productive. The set of these statistics form an extended state space over which dynamic programming may be used to derive the optimal decision rules.

II. Notation and Problem Formulation

A. Notation

This work will use set valued functions

$$\underline{f}: X \rightarrow 2^Y \quad (2.1)$$

to model nondeterministic behavior, i.e. $y \in Y = \underline{f}(x)$ indicates that any element y of the set $Y \subseteq Y$ may arise as a result of applying \underline{f} to the point x . The extension of \underline{f} to a function on the power set of X

$$\underline{f}_e: 2^X \rightarrow 2^Y \quad (2.2)$$

¹Supported by the Office of Naval Research, under contract No.00014-77-0532 (N041-519) is gratefully acknowledged.

²Laboratory of Information and Decision Systems, MIT, Rm. 35-213, Cambridge, MA 02139.

where

$$\underline{f}_e(x) = \bigcup_{x \in x} \underline{f}(x) \quad (2.3)$$

will not be distinguished from \underline{f} itself. The preimage of y under \underline{f} will be denoted by

$$\underline{x} = \underline{f}^{-1}(y) = \{x | y \in \underline{f}(x)\} \quad (2.4)$$

\underline{f}^{-1} will not be distinguished from its extension to

$$\underline{f}_e^{-1}(y) = \{x | y \in \bigcup_{x \in x} \underline{f}(x) \neq \emptyset\} \quad (2.5)$$

Finally, an cost functional

$$J: X \rightarrow \mathbb{R} \quad (2.6)$$

will have an extension

$$J_e(x) = \max_{x \in x} J(x) \quad (2.7)$$

which will also not be distinguished from its restriction, J .

Subscripts will indicate the agent associated with each variable. Superscripts will indicate elements of Cartesian product set; e.g. $x^t = (x(1), \dots, x(t)) \in X^t$.

B. Dynamics

The system to be controlled will have a finite state space X with N elements and dynamics

$$x(t+1) \in \underline{f}(x(t), u_1(t), u_2(t)) \quad (2.8)$$

The state and its dynamics may be taken to include any interagent communication mechanisms. The initial state is assumed fixed, known to all agents and is denoted x_0 .

Observations

$$y_i(t) \in \underline{h}_i(x(t)) \quad (2.9)$$

are available to each agent at each time just before $u_i(t)$ is to be selected.

Each agent is assumed to have perfect recall of all past observations and decisions. The decision rule by which $u_i(t)$ is selected is restricted to being a causal function of the local information sequence

$$Y_i(\cdot; t) : Y_1^t \rightarrow U_i \quad (2.10)$$

Γ_i^T represents the entire sequence of decision rules for agent i .

C. Objective

A cost function is defined as

$$J: X \times U_1 \times U_2 \rightarrow \mathbb{R} \quad (2.11)$$

Taken together, the dynamics (2.8-2.10) and control laws Γ_1^T, Γ_2^T recursively define a set of possible joint state, control, and information trajectories.

ifine

$$J(\Gamma_1^T, \Gamma_2^T) = \max_{t \in \{0, \dots, T\}} \max_{\substack{x(t) \\ u_1(t) \\ u_2(t)}} J(x(t), u_1(t), u_2(t)) \quad (2.12)$$

here the $x(t), u_1(t), u_2(t)$ range over all values jointly possible at time t as determined by this recursion. The overall objective is to minimize this function.

II. The Information Relation

Definition

Begin by considering the autonomous case.

Definition: The global conditional state set at time t

denoted $\hat{x}(t) \subseteq X$, is the set of all possible states which the system may occupy at time t and which may be reached along a trajectory x^* which is possible given both agents' observation sequences.

This global conditional state set is analogous to the Markov conditional state distribution; it may be computed recursively.

Lemma 1: The global conditional state set may be computed from

$$\begin{aligned} \hat{x}(0) &= \{x_0\} \\ \hat{x}(t+1) &= f(\hat{x}(t)) \cap h_1^{-1}(y_1(t+1)) \cap h_2^{-1}(y_2(t+1)) \end{aligned} \quad (3.1)$$

Proof: Set manipulations.

IV. Information Relations

Concept

Definition: An information relation R for a two agent problem is a function from two sets $Z_1(t)$ and $Z_2(t)$ to the power set of another set X . For autonomous systems, R is a primitive information relation if

$$\begin{aligned} Z_1(t) &= Y_1^t, Z_2(t) = Y_2^t, X \text{ is the state space, and} \\ R(Y_1^t, Y_2^t) &= \hat{x}(t). \end{aligned} \quad (4.1)$$

The motivation for studying the information relation is to reduce a primitive relation to one of smaller dimension which still serves for the generation of optimal decision rules. The generic sets Z_i which comprise any relation will be aggregations of the primitive information sets Y_i^t . The remainder of this section establishes the algebraic structure of information relations.

B. Homomorphisms

The global conditional state sets have a lattice structure superimposed on them by the set containment relation. Information relations have a similar structure which, while not a lattice, is (almost) a partial order.

Definition: A homomorphism ϕ from R , defined from Z_1 and Z_2 to 2^X , to R' , defined from Z_1' and Z_2' to 2^X , is a pair of functions $\phi = (\phi_1, \phi_2)$

$$\begin{aligned} \phi_1: Z_1 &\rightarrow Z_1' \\ \phi_2: Z_2 &\rightarrow Z_2' \end{aligned} \quad (4.2)$$

satisfying

$$\begin{aligned} R(z_1, z_2) &\subseteq R'(\phi_1(z_1), \phi_2(z_2)) \quad \forall z_2 \in Z_2 \\ R(\tilde{z}_1, z_2) &\subseteq R'(\phi_1(\tilde{z}_1), \phi_2(z_2)) \quad \forall \tilde{z}_1 \in Z_1 \end{aligned} \quad (4.3)$$

Definition: An information relation R is contained in another R' ($R \subseteq R'$) if there is a homomorphism from R to R' .

The containment relation is clearly reflexive (if ϕ_i is the identity map) and transitive (through composition of ϕ_i 's). It thus imposes a structure on the set of information relations which is partial order on equivalence classes, where:

Definition: R is equivalent to R' ($R \sim R'$) if $R \subseteq R'$ and $R' \subseteq R$.

Both the partial order and equivalence relation have useful interpretations in the decentralized problem.

C. Automorphisms

This section focusses on the equivalence \sim on information relations.

Definition: An automorphism ϕ on an information relation R is a homomorphism from R to R . If both ϕ_1 and ϕ_2 are 1:1, then it is an isomorphism.

Otherwise, it is a reducing automorphism.

The automorphisms which are not isomorphisms are of considerable interest.

Definition: An information relation R is irreducible if all automorphisms on R are isomorphisms. Otherwise it is reducible.

Definition: The reduction of an information relation

$R: Z_1 \times Z_2 \rightarrow 2^X$ by a reducing automorphism

$\phi = (\phi_1, \phi_2)$ is an information relation

$R' : \phi_1(Z_1) \times \phi_2(Z_2) \rightarrow 2^X$ where

$$R'(\phi_1(z_1), \phi_2(z_2)) = R(\phi_1(z_1), \phi_2(z_2)) \quad (4.4)$$

D. The Core

Every information relation can be reduced to an irreducible one by suitable compositions of reducing automorphisms (thus by some single reducing automorphism).

This notion is the most essential part of this work.

Definition: A core of an information relation R , denoted $\text{core}(R)$ or R^* , is an irreducible information relation obtained from R by some automorphism.

The core has a number of interesting properties. The most basic rely on the following lemma.

Lemma 2: If two irreducible information relations R_1 and R_2 are equivalent under \sim , then they are equal.

Proof: If $R \subseteq R'$, then there is a homomorphism ϕ from R to R' , and another, ϕ' from R' to R .

Consider the composed homomorphism $\tilde{\phi}$ from R to R :

$$\begin{aligned} \tilde{\phi}_1(z_1) &= \phi_1'(\phi_1(z_1)) \\ \tilde{\phi}_2(z_2) &= \phi_2'(\phi_2(z_2)) \end{aligned} \quad (4.5)$$

That $\tilde{\phi}$ is indeed a homomorphism from R to R is shown by

$$R(z_1, \tilde{z}_2) \subseteq R'(\phi_1(z_1), \phi_2(\tilde{z}_2)) \quad \forall \tilde{z}_2 \in Z_2 \quad (4.6)$$

since ϕ is a homomorphism from R to R' , and in turn

$$R'(\phi_1(\tilde{z}_1), \phi_2(z_2)) \subseteq R(\phi_1'(\phi_1(\tilde{z}_1)), \phi_2'(\phi_2(z_2))) \quad \forall \tilde{z}_1 \in Z_1 \quad (4.7)$$

¹ Up to isomorphism. This qualifier will be left implicit in the sequel.

by ϕ' being a homomorphism. Thus ϕ is an automorphism on R , and in fact must be an isomorphism since R is irreducible. However, the composition of ϕ and ϕ' can be an isomorphism if and only if ϕ and ϕ' are isomorphisms; hence R and R' are equal.

This lemma immediately gives:

Theorem 1: The core of an information relation is unique. Moreover, if $R \subseteq R'$, then $\text{core}(R) = \text{core}(R')$.

Proof: Any core $(R) \subseteq R$, since a homomorphism from R to a core (R) exists by definition, and a homomorphism from core (R) to R exists by construction: core (R) is a (perhaps relabeled) restriction of R to subsets of Z_1 ; construct the homomorphism from elements of these subsets back into their original values in Z_1 and Z_2 . If two cores, core (R) and core $'(R)$ exist, then

$$\text{core}(R) \subseteq R \subseteq \text{core}'(R) \quad (1.8)$$

Both are irreducible; by transitivity of \subseteq and lemma 2 they must be equal, hence unique.

Moreover

$$\text{core}(R) \subseteq R \subseteq R' \subseteq \text{core}'(R') \quad (4.9)$$

similarly implies $\text{core}(R) = (R')$.

At any point in time $Z_1(t)$ in $R(t)$ will represent a reduced, perhaps trivially, version of Y_1^t . There is a natural way to extend this aggregation to its counterpart at time $t+1$.

Definition: The expansion of an information relation $R(t)$ to another relation $R(t+1)$ is denoted

$$R(t+1) = F(R(t)) \quad (4.10)$$

and is constructed by setting

$$Z_1(t+1) = Z_1(t) \times Y_1 \quad (4.11)$$

and

$$R(t+1)[z_1(t+1), z_2(t+1)] = R(t+1)[z_1(t), y_1, z_2(t), y_2] \quad (4.12)$$

$$\Delta = \underline{f}(R(t)[z_1(t), z_2(t)]) \cap h_1^{-1}(y_1) \cap h_2^{-1}(y_2)$$

The structure of this expansion is captured in

Lemma 3: If $R(t)$ is the primitive relation at time t , then $F(R(t))$ is the primitive relation for time $t+1$.

Proof: By definition, $R(t): Y_1^t \times Y_2^t \rightarrow Z^X$, and

$R(t)(y_1^t, y_2^t) = \hat{x}(t)$, the global conditional state set based on (y_1^t, y_2^t) . Identifying y_1 in (4.13) with the observation $y_1(t+1)$ yields

$$\begin{aligned} R(t+1)(y_1^{t+1}, y_2^{t+1}) &= \\ &= \underline{f}(\hat{x}(t)) \cap h_1^{-1}(y_1(t+1)) \cap h_2^{-1}(y_2(t+1)) \quad (4.13) \\ &= \hat{x}(t+1) \end{aligned}$$

by (4.1). Thus $R(t+1)$ is a primitive information relation.

Lemma 4: If $R \subseteq R'$, then

$$F(R) \subseteq F(R') \quad (4.14)$$

Proof: Since $R \subseteq R'$, there exists functions $\phi_i: Z_1 \rightarrow Z_1'$ satisfying (4.4). Construct functions

$$\phi_i^+: Z_1 \times Y_1 \rightarrow Z_1' \times Y_1 \text{ where}$$

$$\phi_i^+(z_1, y_1) = (\phi_i(z_1), y_1) \quad (4.15)$$

Then for all \tilde{z}_2 and \tilde{y}_2 ,

$$F(R((z_1, y_1), \tilde{z}_2, \tilde{y}_2))) = \underline{f}(R(z_1, \tilde{z}_2)) \cap h_1^{-1}(y_1) \cap h_2^{-1}(\tilde{y}_2) \quad (4.16)$$

$$= \underline{f}(R'(\phi_1(z_1), \phi_2(\tilde{z}_2))) \cap h_1^{-1}(y_1) \cap h_2^{-1}(\tilde{y}_2) \quad (4.17)$$

by (4.4) and set inequalities, and this

$$= F(R')(\phi_1^+(z_1, y_1), \phi_2^+(\tilde{z}_2, \tilde{y}_2)) \quad (4.18)$$

This establishes the first half of (4.4) for ϕ^+ ; the other half is shown by a symmetric argument. Thus ϕ^+ is a homomorphism from $F(R)$ to $F(R')$; hence $F(R) \subseteq F(R')$.

This sets up the second major result:

Theorem 2: Let $R(t)$ and $R(t+1)$ be primitive information relations at successive times. Then

$$\text{core}(R(t+1)) = \text{core}(F(\text{core}(R(t)))) \quad (4.19)$$

Proof: By Lemma 3.

$$R(t+1) = F(R(t)) \quad (4.20)$$

From Theorem 1

$$\text{core}(R(t)) \subseteq R(t) \quad (4.21)$$

Hence applying lemma 4 to each containment in (4.21)

$$F(\text{core}(R(t))) \subseteq F(R(t)) \quad (4.22)$$

Then by Theorem 1 and (3.20)

$$\text{core}(F(\text{core}(R(t)))) = \text{core}(R(t+1)) \quad (4.23)$$

Definition: The steady state core information relation R^* , if it exists, is the core of some primitive information relation $R(t)$ and satisfies

$$R^* = \text{core}(F(R^*)) \quad (4.24)$$

If a steady state core can be found for a system, then a great deal of the work required to solve the system is complete. However, not all systems have a steady state core.

The core dynamics iteratively expand $Z_1(t)$ by appending Y_1 , a new observation, and then reduce it via an equivalence relation implied by an automorphism ϕ_i . The equivalence relation combines those elements of $Z_1 \times Y_1$ which need not be distinguished in the future as far as the core dynamics are concerned. The entire sequence of these equivalence relations map every observation sequence y_1^t into some element of $Z_1(t)$, and thus dictate the structure of a finite state machine with Y_1 as an input set and $Z_1(t)$ as the states.

Definition: The local core observer for agent i is the system with state set $Z_1(t)$ and dynamics

$$\begin{aligned} z_1(t+1) &= \hat{f}_i(z_1(t), y_1(t+1)) \\ \Delta &= \phi_i(z_1(t), y_1(t+1)) \end{aligned} \quad (4.25)$$

where the $Z_1(t)$ are the components on which $R^*(t)$ is defined, and ϕ_i is a component of the reducing automorphism used to reduce $F(R^*(t))$ to $R^*(t+1)$.

Thus the core dynamics define some automata for processing local observations in a way which maintains the relationship described in R^* .

V. Decentralized Estimation

A. Problem

The general decentralized estimation problem is to find the

$$\min_{\Gamma_1^T, \Gamma_2^T} \max_{t \in 1, \dots, T} \max_{x(t) \text{ possible}} J(x(t), u_1(t), u_2(t)) \quad (5.1)$$

where $x(t)$ is possible if $x(t) \in \underline{f}^t(x_0)$, (the t -fold composition of \underline{f} with itself). The information restriction

$$u_i(t) = \gamma_i(t) \quad (\gamma_i^t) \quad (5.2)$$

still applies.

Define

$$\bar{J}(\gamma_1(t), \gamma_2(t)) = \max_{\substack{x(t) \\ \gamma_1 \\ \gamma_2 \\ t}} J(x(t), \gamma_1(x_1^t), \gamma_2(y_2^t)) \quad (5.3)$$

possible

Each component of Γ_i^T may be chosen separately, since decisions do not affect dynamics or costs other than that a single time; J depends only on decision rules at one time. Thus

Lemma 5: The solution to the decentralized estimation problem may be found by solving the sequence:

$$\text{minimize } \bar{J}(\gamma_1(t), \gamma_2(t)) \quad (5.4)$$

$\gamma_1(t), \gamma_2(t)$

B. Use of the Information Relation

Definition: The optimal value of an information relation R in a nondeterministic decentralized estimation problem, denoted $J^*(R)$, is

$$\min_{\gamma_1, \gamma_2} \max_{\substack{z_1 \in Z_1 \\ z_2 \in Z_2 \\ x \in R(z_1, z_2)}} J(x, \gamma_1(z_1), \gamma_2(z_2)) \quad (5.5)$$

with Z_1 and Z_2 the sets on which R is defined, and the restriction

$$\gamma_i : Z_i \rightarrow U_i \quad (5.6)$$

Lemma 6: If R and R' are information relations, then

$$R \subseteq R' \Rightarrow J^*(R) \leq J^*(R') \quad (5.7)$$

Proof: If $R \subseteq R'$, then a homomorphism $\phi = (\phi_1, \phi_2)$

exists from R to R' . Consider any strategy

$\gamma' = (\gamma'_1, \gamma'_2)$ with

$$\gamma'_i : Z'_i \rightarrow U_i \quad (5.8)$$

Build a strategy $\gamma = (\gamma_1, \gamma_2)$ with

$$\gamma_i(z_i) = \gamma'_i(\phi_i(z_i)) \quad (5.9)$$

Then the inequality stems from the fact that

$$\phi_i^* Z'_i \subseteq Z_i \quad (5.10)$$

$$R(z_1, z_2) \subseteq R'(\phi_1(z_1), \phi_2(z_2))$$

from the definition of a homomorphism. The γ defined in (4.9), must achieve a value no larger than $J^*(R')$, and the best γ is at least as good as this one, so the conclusion holds. This immediately suggests the third major result:

Theorem 3: Let $R(t)$ be a primitive information relation and $R^*(t)$ its core.

Then

$$J^*(R(t)) = J^*(R^*(t)) \quad (5.11)$$

Proof: From theorem 1, $R(t) \subseteq R^*$. Applying Lemma 6 in both directions, the conclusion follows.

VI. The Control Problem.

A. Information Dynamics

The definitions of information relation, homomorphisms, containment, and cores established in section IV carry over to the control problem without change. The expansion process is the only place where new information relations are generated, so the influence of decisions on dynamics requires a modification there.

Definition: The expansion of an information relation R by decision rules $\gamma_i : Z_i \rightarrow U_i$ is an information relation R' , where

$$R' = F(R, \gamma_1, \gamma_2) \quad (6.1)$$

$$R' : Z'_1 \times Z'_2 \rightarrow Z^X \quad (6.2)$$

$$Z'_i = Z_i \times Y_i \quad (6.3)$$

$$R'((z_1, y_1), (z_2, y_2)) = \underline{f}(R(z_1, z_2), \gamma_1(z_1) \gamma_2(z_2)) \quad (6.4)$$

$\bigcap_{h=1}^{-1} (y_1) \bigcap_{h=2}^{-1} (y_2)$

Definition: Let ϕ be a homomorphism from R to R' .

Then a decision rule $\gamma_i : Z_i \rightarrow U_i$ is contained in

$$\gamma'_i : Z'_i \rightarrow U_i, \text{ denoted } \gamma_i \subseteq \gamma'_i,$$

if

$$\gamma_i(z_i) = \gamma'_i(\phi_i(z_i)) \quad \forall z_i \in Z_i \quad (6.5)$$

With this notion, the results of section IV generalize to the control case as:

Theorem 4: Let R and R' be information relations, and γ_1, γ_2 decision rules of the appropriate structure.

a) if R is a primitive information relation, then so is $F(R, \gamma_1, \gamma_2)$

b) if $R \subseteq R'$, $\gamma_1 \subseteq \gamma'_1$ and $\gamma_2 \subseteq \gamma'_2$, all by the same homomorphism, then

$$F(R, \gamma_1, \gamma_2) \subseteq F(R', \gamma'_1, \gamma'_2) \quad (6.6)$$

c) If R is a primitive information relation, then $\text{core}(R(\gamma_1, \gamma_2)) = \text{core}(F(\text{core}(R), \gamma_1, \gamma_2))$ (6.7)

Proof: All proofs are direct extension of

a) Lemma 3

b) Lemma 4

c) Theorem 2, using (a) and (b).

B. Costs

Properties developed in Section IV for the estimation problem also generalize to the control problem.

Definition: The cost of an information relation R with compatible decision rules γ_1, γ_2 is

$$J(R, \gamma_1, \gamma_2) = \max_{\substack{z_1 \\ z_2 \\ x \in R(\gamma_1, z_2)}} J(x, \gamma_1(z_1), \gamma_2(z_2)) \quad (6.8)$$

The overall problem objective (2.13) becomes

$$J(\Gamma_1, \Gamma_2) = \max_{t \in \{1, \dots, T\}} \max_{\substack{y_1^t \\ y_2^t \\ x \in R(y_1^t, y_2^t)}} J(x(t), \gamma_1(y_1^t), \gamma_2(y_2^t)) \quad (6.9)$$

$$= \max_{t \in \{1, \dots, T\}} J(R(t), \gamma_1(t), \gamma_2(t)) \quad (6.10)$$

where the dependence of $R(t)$ on prior decision rules is left implicit. Thus the overall objective can be written in terms of costs of primitive relations with decision rules.

The final result needed is:

Theorem 5: Let R and R' be information relations, with γ_1, γ_2 and γ'_1, γ'_2 compatible decision rules.

Then

- a) If $R \subseteq R'$, $\gamma_1 \subseteq \gamma'_1$, and $\gamma_2 \subseteq \gamma'_2$ by the same homomorphism ϕ , then

$$J(R, \gamma_1, \gamma_2) \leq J(R', \gamma'_1, \gamma'_2) \quad (6.11)$$

- b) If $R^* = \text{core}(R')$, $\gamma_1 \subseteq \gamma'_1$, and $\gamma_2 \subseteq \gamma'_2$ by the same homomorphism ϕ , then

$$J(R^*, \gamma_1, \gamma_2) = J(R', \gamma'_1, \gamma'_2) \quad (6.12)$$

Proof: Also direct extensions of previous results.

- a) Lemma 6
b) (a) with Theorem 1.

C. General Solution

Let $R^*(t) = \text{core}(R(t))$, where $R(t)$ is the primitive information relation created by decision rules prior to time t . From Theorem 4c,

$$\begin{aligned} R^*(t+1) &= \text{core}(F(R^*(t), \gamma_1(t), \gamma_2(t))) \\ &\stackrel{\Delta}{=} F^*(R^*(t), \gamma_1(t), \gamma_2(t)) \end{aligned} \quad (6.13)$$

and the overall objective is

$$J(\Gamma_1, \Gamma_2) = \max_{t \in \{1, \dots, T\}} J(R^*(t), \gamma_1(t), \gamma_2(t)) \quad (6.14)$$

where Γ_1 and Γ_2 are now sequences of decision rules whose arguments are in $Z_1(t)$ and $Z_2(t)$ respectively.

Solution of this problem would involve straightforward minimax dynamic programming if the set in which the $R^*(t)$ could like were determinable ahead of time. Unfortunately, it is not known how to do this at this point; however, one may construct the set of all cores reachable under all Γ_1, Γ_2 from $R^*(0)$ and take this to be the requisite set.

A MINIMUM SENSITIVITY INCENTIVE CONTROL APPROACH TO TEAM PROBLEMS*

Tamer Başar, Jose B. Cruz, Jr., and Derya Cansever
 Department of Electrical Engineering and
 Coordinated Science Laboratory
 University of Illinois
 1101 W. Springfield Ave.
 Urbana, Illinois 61801

ABSTRACT

In this paper we analyze a class of two-agent team decision problems with a hierarchical decision structure, wherein one of the decision makers may have a slightly different perception of the overall team goal, with this slight variation not known by the other agent who is assumed to occupy the hierarchically dominant position. The leading agent has access to dynamic information and his role is to announce such a policy (incentive scheme) which would lead to achievement of the overall team goal, in spite of the slight variations in the other agent's perception of that goal, which are not known or predictable by him. We may call a policy with such an additional feature a "minimum sensitivity" incentive policy. We obtain, in the paper, "minimum sensitivity" policies for the leading agent, for a general cost functional with convex structure, which are least sensitive to variations in the following agent's perception of the team goal. In some special cases, we show that the robust feature of the incentive scheme is maintained regardless of the magnitude and nature of the variations, and illustrate the theory with an example arising in armament limitation and control.

I. INTRODUCTION

The main characteristic of team decision problems is the presence of several decision makers with a common objective functional which is to be optimized jointly (but possibly in a decentralized fashion) by all decision makers. An underlying stipulation in research on team theory has been the assumption that all agents perceive the common goal in exactly the same way, and face exactly the same mathematical optimization problem [Marschak and Radner (1972)]. In this paper we relax this basic assumption and allow (in the context of two-agent problems) one agent to have a somewhat different perception of the common goal and to quantify it in a slightly different way. Furthermore, we will assume that the other agent is not informed of the existence of this discrepancy in the perception of the common goal, but is able to monitor the decision of the former by occupying a higher (dominant) position in the decision process. The problem we address to is the design of a suitable strategy for the agent who occupies the hierarchically superior position and who still adopts the original team objective functional as his own, such that the change in the minimum value of the team cost because of the discrepancy in the perceptions of the common goal is kept to a minimum. Ideally, the hierarchically superior member of the team would seek not to be affected by this discrepancy, if this is at all possible.

We will approach this problem using optimum incentive design schemes [Ho, Luh, and Olsder (1982) and

Zheng and Başar (1982)], which involve a hierarchy in decision making and a suitable information structure for the decision maker at the top of the hierarchy, that allows him to design a policy which in its turn induces the other decision maker with a different objective functional to behave in a desired manner. Recently in [Cansever and Başar (1982)], optimal incentive schemes have been used, within the context of Stackelberg games, to minimize the effect of changes in the parameters of the follower's cost functional on the leader's optimum cost value, by simultaneously achieving a desired goal. Here, we direct our attention to problems which are nominally team, and derive incentive schemes that are least sensitive to deviations in the hierarchically inferior decision maker's perceptions of the uncertain parameters. The fact that the underlying goal is common (that is, the nominal optimization problem is a team problem—a property that may be destroyed in the decision process) can be exploited to obtain very appealing minimum sensitivity strategies, as we will show in the sections to follow.

The problem is formulated in Section II. In Section III we introduce sensitivity functions and obtain robust affine strategies for a general class of convex cost functionals. In Section IV we provide a geometrical interpretation for total insensitivity when the objective functional is affine in the unknown parameter. Section V deals with the generalization of some of these results to the multiparameter case, and Section VI illustrates the basic ideas developed in this paper using a model from armament limitation and control. Concluding remarks of Section VII end the paper.

II. PROBLEM FORMULATION

Consider a two-person deterministic team decision problem in normal form, described by the cost functional $J(y_1, y_2, u)$, where $y_i \in \Gamma_i$ denotes the strategy of DMi (i'th decision maker) and $u \in ACR$ is a parameter on which the cost functional depends. Let $u \in U = R^n$, $v \in V = R^m$ denote the decision variables of DM1 and DM2, respectively, and assume that $\Gamma_1 = \{y: V \rightarrow U\}$, $\Gamma_2 = V$; i.e. DM1 has access to the decision value of DM2. DM1 also knows the precise value of the parameter u (say u^0), whereas DM2 perceives its value differently (say $u^+ \in A$), which in turn gives rise to a different cost functional from his point of view, namely, $J(y_1, y_2, u^+) \neq J(y_1, y_2, u^0)$. Furthermore, DM1 does not know the exact value perceived by DM2, but his ultimate goal is to see that the lowest possible value is attained for $J(y_1, y_2, u^0)$. The decision structure of the problem is assumed to be hierarchical, in the sense that DM1 is the dominant decision maker and has the power and means of declaring his policy in advance and enforcing it on the other DM. Hence, while DM2 is faced with the problem of minimizing $J(y_1(v), v, u^+)$ over $v \in V$, DM1 wishes to choose a $y_1 \in \Gamma_1$ (in total ignorance of u^+) that would eventually lead to a minimum value for $J(y_1(v), v, u^0)$.

By an abuse of notation, let $J(u, v, a)$ denote the cost functional on the product space $U \times V$, for each

*This work was supported by the Office of Naval Research under Contract N00014-82-K-0469.

$\alpha \in A$, and assume that this functional is strictly convex on $U \times V$, for each $\alpha \in A$, is twice continuously differentiable in its first two arguments and continuously differentiable in its third argument. Furthermore, let us denote the unique minimum of $J(u, v, \alpha^0)$ by $(u^t, v^t) \in U \times V$. Restricting DM1 to affine policies in Γ_1 , we first note that the policy

$$\gamma_1(v) = u^t + P(v - v^t), \quad (1)$$

where P is an $(n \times m)$ -matrix, has the appealing property that if DM2's perception of α is α^0 , then $\min_{v \in V} J(\gamma_1(v), v, \alpha^0)$ leads to the desired value $v^t \in V$ for any matrix P . If $\alpha^+ \neq \alpha^0$, however, the problem ceases to be a cooperative one since the problem faced by DM2 is

$$\min_{v \in V} J(\gamma_1(v), v, \alpha^+) \quad (2a)$$

whose minimizing solution (say $v^+ \in V$) satisfies (and is uniquely determined by) the equation

$$J_u(u^+, v^+, \alpha^+)P + J_v(u^+, v^+, \alpha^+) = 0 \quad (2b)^1$$

where $u^+ = \gamma_1(v^+) = u^t + P(v^+ - v^t)$ and is not necessarily the same as u^t . The problem we address, in the sequel, is whether it is possible to choose a robust policy γ_1 (by choosing P appropriately) so that either $u^+ = u^t$ and $v^+ = v^t$, or the discrepancies will be small whenever α^+ is close to α^0 ; in other words, we seek either total insensitivity or minimum sensitivity of the optimum value of $J(u, v, \alpha^0)$ to variations in the perception of DM2 (of α) by a proper choice of γ_1 .

III. INTRODUCTION OF A SENSITIVITY FUNCTION AND DERIVATION OF MINIMUM SENSITIVITY SOLUTIONS

As a measure of the sensitivity of $J(u, v, \alpha^0)$ with respect to deviations in the perception of DM2 of α from its nominal value α^0 , let us introduce the total derivative of $J(u, v, \alpha^0)$ with respect to α , when $u = u^t$, $v = v^t$, satisfying (2b), and at the point $\alpha^+ = \alpha^0$. We call this function the *first-order sensitivity function* of $J(u, v, \alpha^0)$ with respect to α , at $\alpha = \alpha^0$, in view of (1) and the optimal response of DM2 as characterized (uniquely) by (2b):

$$\begin{aligned} I_1(\alpha^0) &= dJ(u^+, v^+, \alpha^0)/d\alpha^+ \Big|_{\alpha^+ = \alpha^0} \\ &= (dJ/dv^+) (dv^+/d\alpha^+) \Big|_{\alpha^+ = \alpha^0} \\ &= [J_u(u^t, v^t, \alpha^0)P + J_v(u^t, v^t, \alpha^0)]v_\alpha^t, \end{aligned} \quad (3)$$

where

$$v_\alpha^t = dv^+(\alpha^+) / d\alpha^+ \Big|_{\alpha^+ = \alpha^0}$$

and is determined from (2b). To obtain an expression for v_α^t , we note that (2b) is in fact an identity for all $\alpha^+ \in A$, since it uniquely determines the optimal response of DM2 to the announced policy (1) of DM1, with his perceived value for α being α^+ . Hence, differentiating (2b) with respect to α^+ , and evaluating the resulting expression at $\alpha^+ = \alpha^0$, we obtain

$$[P'J_{uu}P + P'J_{vu} + J_{uv}P + J_{vv}]v_\alpha^t + [J_{au}P + J_{av}] = 0 \quad (4)$$

whereby

$$v_\alpha^t = -[P'J_{uu}P + P'J_{vu} + J_{uv}P + J_{vv}]^{-1} [J_{au}P + J_{av}] \quad (5)$$

¹⁾ Here J_u and J_v are row vectors of dimensions $1 \times n$ and $1 \times m$, respectively, denoting the partial derivatives with respect to the corresponding decision variables.

where the arguments are evaluated at $u = u^t$, $v = v^t$, $\alpha = \alpha^0$. Note that the required inverse in (5) exists under the initial hypothesis that J is strictly convex in (u, v) for all $\alpha \in A$.

Now, since the pair (u^t, v^t) globally minimizes $J(u, v, \alpha^0)$, we already know that

$$J_u(u^t, v^t, \alpha^0) = 0, \quad J_v(u^t, v^t, \alpha^0) = 0, \quad (6)$$

in view of which the first product term of (3) and hence $I_1(\alpha^0)$ vanish. Then, the dominating term in the Taylor expansion of $J(u^+, v^+, \alpha^0)$ around $\alpha^+ = \alpha^0$ is determined by the second-order sensitivity function:

$$\begin{aligned} I_2(\alpha^0) &= d^2J(u^+, v^+, \alpha^0)/d\alpha^{+2} \Big|_{\alpha^+ = \alpha^0} \\ &= [(dv^+/d\alpha^+)(d^2J/dv^{+2})(dv^+/d\alpha^+) \\ &\quad + (dJ/dv^+)(d^2v^+/d\alpha^{+2})] \Big|_{\alpha^+ = \alpha^0} \\ &= v_\alpha^t [P'J_{uu}P + P'J_{vu} + J_{uv}P + J_{vv}]v_\alpha^t + [J_uP + J_v]v_{\alpha\alpha}^t. \end{aligned} \quad (7)$$

Since the second term is zero, in view of (6), $I_2(\alpha^0)$ vanishes if and only if $v_{\alpha\alpha}^t = 0$, which requires from (5) that there exist a P satisfying

$$J_{au}(u^t, v^t, \alpha^0)P + J_{av}(u^t, v^t, \alpha^0) = 0 \quad (8)$$

A sufficient condition for this is, of course,

$$J_{au}(u^t, v^t, \alpha^0) \neq 0, \quad (9)$$

which is also necessary if the second term in (8) does not vanish (at least one component is nonzero).

When v_α^t vanishes, not only the second-order sensitivity function, but also the third-order sensitivity function

$$I_3(\alpha^0) = d^3J(u^+, v^+, \alpha^0)/d\alpha^{+3} \Big|_{\alpha^+ = \alpha^0} \quad (10)$$

vanishes, because it carries (by chain rule of differentiation) only terms that involve either v_α^t or $[J_uP + J_v]$ as products. Hence, under the condition that (8) admits at least one solution, and when DM1 employs the corresponding policy (1), if DM2's perception α^+ (of α) stays within an ϵ -neighborhood of its nominal value α^0 , the 3rd order Taylor approximation of the effect of this discrepancy is zero. We now summarize this appealing feature of the linear policy (1) in the following proposition.

Proposition 1: Let condition (9) be satisfied, and let P^* denote a solution to (8). Then, if DM1 employs the policy

$$\gamma_1^*(v) = u^t + P^*(v - v^t), \quad (11a)$$

and the unique minimizing decision of DM2 is

$$v^+(\alpha^+) = \arg \min_{v \in R^m} J(\gamma_1^*(v), v, \alpha^+), \quad (11b)$$

$J(\gamma_1^*(v^+), v^+, \alpha^0)$ agrees with $J(u^t, v^t, \alpha^0)$ to third order in α^+ when it lies in a sufficiently small neighborhood of α^0 . Equivalently, the discrepancy in costs is of fourth order.

When the objective function $J(u, v, \alpha)$ is affine in α , we can obtain more explicit results. Specifically, let

$$J(u, v, \alpha) = g(u, v) + \alpha h(u, v) \quad (12a)$$

where g and h are continuously differentiable in their arguments,

$$h_u(u^t, v^t) \neq 0, \quad (12b)$$

and J is strictly convex in (u, v) for all $\alpha \in A$. Then, (8) reads

$$h_u(u^t, v^t)P + h_v(u^t, v^t) = 0, \quad (13)$$

a solution to which always exists because (12b) becomes equivalent to (9). Hence v_α^t , as given by (5) [evaluated at $\alpha = \alpha^0$] is zero. This, in turn, implies through an iterative verification that the vector $d^n v^+(\alpha^+) / d\alpha^+{}^n$, where $v^+(\alpha^+)$ is given by (11b), vanishes at $\alpha^+ = \alpha^0$, for all $n = 1, 2, \dots$, simply because the second term in (4)

$$J_{\alpha u}P + J_{\alpha v} = h_u(u, v)P + h_v(u, v)$$

is not explicitly dependent on α . Since the n th order sensitivity function

$$I^n(\alpha^0) = d^n J(u^+, v^+, \alpha^0) / d\alpha^+{}^n \Big|_{\alpha^+ = \alpha^0} \quad (14)$$

carries only $(d^i v^+(\alpha^+) / d\alpha^+{}^i)_{\alpha^+ = \alpha^0}$, $i = 1, 2, \dots, n$ as product terms, which are all zero whenever P is chosen to satisfy (13), it follows that sensitivity functions of all orders vanish, at $\alpha^+ = \alpha^0$. Hence,

Proposition 2: When the objective function is given by (12a), under the condition (12b), let P^* be any solution of (13). Then, if (11a) is employed by DM1, the response of DM2 (i.e., (11b)) is independent of α^+ , and $v^+ = v^t$. Hence, $J(u^+, v^+, \alpha^0) = J(u^t, v^t, \alpha^0)$ for all $\alpha^+ \in A$, that is the overall performance is independent of the perception of DM2 regarding the value of α . \square

In the next section we provide a geometric interpretation of this appealing feature of the linear policy when the cost function is an affine function of the parameter α .

IV. GEOMETRIC INTERPRETATION OF TOTAL INSENSITIVITY WHEN THE OBJECTIVE FUNCTIONAL IS AFFINE IN A PARAMETER

Let the objective function J be as given by (12a) with h satisfying condition (12b). Since J is strictly convex, the team solution (u^t, v^t) when $\alpha = \alpha^0$ is obtained (uniquely) from

$$g_u(u^t, v^t) + \alpha^0 h_u(u^t, v^t) = 0 \quad (15a)$$

$$g_v(u^t, v^t) + \alpha^0 h_v(u^t, v^t) = 0. \quad (15b)$$

Postmultiplying (15a) by P , adding this to (15b), and taking the transpose, we have

$$(P'g_u' + g_v') + \alpha^0 (P'h_u' + h_v') = 0. \quad (16)$$

Pictorially, the vectors $(P'g_u' + g_v')$ and $(P'h_u' + h_v')$ are oppositely oriented when α^0 is a positive scalar. Clearly, α^0 is the ratio of the magnitude of the vector $(P'g_u' + g_v')$ to the magnitude of the vector $(P'h_u' + h_v')$. If DM1 chooses P such that (13) is satisfied, then the magnitudes of both vectors become zero. In this case, if α^0 is replaced by $\alpha \neq \alpha^0$ in (16), the equation would still hold, and (u^t, v^t) satisfies

$$(P'g_u' + g_v') + \alpha (P'h_u' + h_v') = 0. \quad (17)$$

Since (17) is the condition used by DM2 to optimize v (see also (2b)), he will choose $v = v^t$, no matter what his perceived value of α is. Thus, DM1 achieves the

team-optimal solution for minimizing $J(u, v, \alpha^0)$ by choosing P such that P' transforms the vector h_u' to $(-h_v')$:

$$P'h_u'(u^t, v^t) = -h_v'(u^t, v^t). \quad (18a)$$

This same choice of P' transforms g_u' to $(-g_v')$:

$$P'g_u'(u^t, v^t) = -g_v'(u^t, v^t). \quad (18b)$$

V. EXTENSION TO THE MULTIPARAMETER CASE

In the previous sections, we have restricted our discussion to the case $\alpha \in \mathbb{R}$. When $\alpha \in \mathbb{R}^r$, the first-order sensitivity function $I_1(\alpha^0)$ becomes a $(1 \times r)$ vector given by

$$I_1(\alpha^0) = dJ(u^+, v^+, \alpha^0) / d\alpha^+ \Big|_{\alpha^+ = \alpha^0} = [J_u(u^t, v^t, \alpha^0)P + J_v(u^t, v^t, \alpha^0)]v_\alpha^t \quad (19a)$$

where

$$v_\alpha^t = -[P'J_{uu}P + P'J_{vu} + J_{uv}P + J_{vv}]^{-1}[J_{u\alpha}P + J_{v\alpha}] \quad (19b)$$

and the arguments are evaluated at $u = u^t$, $v = v^t$, $\alpha = \alpha^0$. Note that $I_1(\alpha^0) = 0$, in view of (6). Furthermore, $v_\alpha^t = 0$ (zero matrix) if

$$J_{v\alpha}(u^t, v^t, \alpha^0) \in \mathcal{R}(J_{u\alpha}(u^t, v^t, \alpha^0)) \quad (20)$$

(where \mathcal{R} denotes the range) since then it is possible to find an $(n \times m)$ matrix P to make the second product term of (19b) zero. In this case, sensitivity functions of orders 1, 2, and 3 vanish at the nominal solution point; hence, affine policies have very appealing sensitivity properties also in the multiparameter case. When condition (20) is not satisfied, however, one has to minimize a suitable norm of the leading sensitivity function with respect to the $(n \times m)$ matrix P . This is, in general, the second-order sensitivity functions $I_2(\alpha^0)$ which is an $(r \times r)$ nonnegative definite matrix. A suitable norm for minimization is, in this case, $\text{Tr}(I_2(\alpha^0))$. We are now faced with an unconstrained optimization problem on P , for which a closed-form solution does not in general exist; however, numerically it is a feasible problem.

When the objective function J is affine in the parameter vector $\alpha \in \mathbb{R}^r$, a total insensitivity result could be established under certain conditions, by a direct extension of the discussion of Section IV. Towards this end, let

$$J(u, v, \alpha) = g(u, v) + \alpha'h(u, v) \quad (21)$$

where $g: U \times V \rightarrow \mathbb{R}$, $h: U \times V \rightarrow \mathbb{R}^r$, J is strictly convex and continuously differentiable in (u, v) . Then, the optimality conditions for $\alpha = \alpha^0$ are

$$g_u(u^t, v^t) + \alpha^0 h_u(u^t, v^t) = 0 \quad (22a)$$

$$g_v(u^t, v^t) + \alpha^0 h_v(u^t, v^t) = 0 \quad (22b)$$

where h_u (respectively, h_v) is an $r \times n$ (respectively, $r \times m$) matrix. The optimal response of DM2, under the policy (1) for DM1, is determined uniquely from (for a general α)

$$(P'g_u' + g_v') + \sum_{i=1}^r \alpha_i (P'h_{iu}' + h_{iv}') = 0 \quad (23)$$

where subscript i denotes the i 'th component of the corresponding vector. Now, let us assume that there exists an $m \times n$ matrix P satisfying simultaneously

$$h_{iu}(u^t, v^t)P + h_{iv}(u^t, v^t) = 0, \quad i=1, \dots, r. \quad (24)$$

Under this condition, the second term in (23) vanishes at $v=v^t$, for all $\alpha \in AC \mathbb{R}^r$, and furthermore the first term also vanishes in view of (22a)-(22b), by basically following the argument of Section IV. Hence, under this particular choice of P , $v=v^t$ is the unique solution to (23) for all values of α ; that is, the optimal response of DM2 is independent of his perception of the value of α , provided that strict convexity of J is preserved. To summarize,

Proposition 3: When the objective function is given by (21), let there exist a solution to (24), to be denoted P^* . Then, if the policy

$$v_1^*(v) = u^t + P^*(v - v^t)$$

is employed by DM1, the response of DM2 (which is (11b) with $\alpha^+ \in \mathbb{R}^r$) is independent of α^+ , and $v^+ = v^t$. Consequently, $J(u^t, v^t, \alpha^0) = J(u^t, v^t, \alpha^0)$ for all $\alpha^+ \in AC \mathbb{R}^r$. \square

In the next section we consider an example that involves arms race between two nations, which serves to illustrate some of the ideas generated in this and the previous sections. Another example from microeconomics can be found in [Cansever, Başar, and Cruz (1983)].

VI. AN EXAMPLE FROM THE PROBLEM OF ARMAMENT LIMITATION AND CONTROL

In their papers on armament race and control [Simaan and Cruz (1975a) and Simaan and Cruz (1975b)], Simaan and Cruz have modeled the arms race problem as a noncooperative differential game between two nations. A salient feature of this model is that, when the respective cost functionals are taken to be quadratic in the decision variables, the resulting optimal state trajectory yields a discretized version of the armament model proposed earlier by Richardson [Richardson (1960)]. We will consider here the case when the two nations, DM1 and DM2, have agreed to reduce their respective armament expenditures. Such a situation inevitably requires the presence of an element of cooperation between DM1 and DM2, since any significant departure from the armament level jointly agreed upon may eventually lead to the original high armament expenditure. Towards the formulation of this problem, let us assume that the goals of the DM's can be represented by two objective functionals $J_i(x_1, x_2, u_1, u_2)$, $i=1,2$, wherein DM1 aims to optimize J_1 . In order to incorporate the cooperation element discussed above, we will adopt the Pareto optimal equilibrium concept, which will be realized [Schmitendorf and Leitmann (1974)] if the DM's jointly optimize

$$J(x_1, x_2, u_1, u_2) = k_1 J_1(x_1, x_2, u_1, u_2) + k_2 J_2(x_1, x_2, u_1, u_2) \quad (25)$$

where $k_i \in \mathbb{R}_+$; $u \in \mathbb{R}_+$, and $u_2 \in \mathbb{R}_+$ denote DM1 and DM2's²⁾ armament investments, respectively, and x_i represents the armament level of DM1, $i=1,2$, which further satisfies

$$x_i = f_i(x_{i0}, u_i), \quad i=1,2. \quad (26)$$

²⁾ \mathbb{R}_+ denotes the positive real line.

where $f_i: \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a continuous function of x_{i0} and u_i , and is strictly increasing in its second argument. Here, x_{i0} denotes the initial armament level of DM1.

In order to obtain some explicit results, let us adopt the quadratic objective functional model proposed by Simaan and Cruz [Simaan and Cruz (1975a) and Simaan and Cruz (1975b)], because of its analytical tractability and other appealing features in relation with other existing models; namely, let

$$J_i(x_1, x_2, u_1, u_2) = \frac{1}{2} \{R_i(u_i - w_i)^2 + Q_i^0(x_i - S_i x_j - v_i)^2\} \quad (27)$$

and

$$x_i = f_i(x_{i0}, u_i) = \delta_i x_{i0} + u_i, \quad i=1,2; j=1,2; i \neq j \quad (28)$$

where

$$R_i > 0, \quad Q_i^0 \geq 0, \quad S_i \geq 0, \quad 0 < \delta_i \leq 1, \quad i=1,2.$$

Here, x_{i0} denotes the given initial armament level of DM1 and expression (27) reveals the fact that each DM wants to reduce the gap that exists between his armament level and a linear function of the other DM's armament level, and at the same time wishes to minimize his expenditure. We refer to [Simaan and Cruz (1975b)] for an elaborated interpretation of (27). Under this set up, there exists a unique pair (u_1^t, u_2^t) , minimizing $J(x_1, x_2, u_1, u_2)$ as a function of $\frac{k_2 \Delta}{k_1} \alpha^0$, which corresponds to the pair (u^t, v^t) in the general discussion of Sections II and V.

As it may be the case, one of the DM's, say DM2, may deviate from u_2^t . The reason behind such a move may be that DM2 totally ignores the cooperation, and minimizes his own objective functional. Assuming that each DM can monitor the decisions of his adversary, this situation would immediately give rise to a Nash equilibrium with high armament expenditures. Since we have assumed that each DM desires to reduce his expenditures while maintaining a certain balance of powers, such a unilateral and large deviation will be unlikely. In its stead, we will assume that DM2 may have an incentive to perform a relatively small deviation from the Pareto equilibrium point, being motivated by one of the following three considerations:

i) DM2 may decide to promote his relative importance in the agreement, which is reflected by an increase in the value of α from α^0 to α^+ , without informing DM1, while DM1 still uses the value α^0 in his objective functional;

ii) DM2 may develop a different perception of the values of one or more coefficients in the team objective functional without informing DM1. Let us assume, for instance, that DM2 has decided to place higher priority and emphasis on reducing the gap between his armament level and the linear functional of DM1's armament level than on minimizing his expenditure; more precisely, that he has decided to increase the value of Q_2 from Q_2^0 to Q_2^+ .

iii) Both i) and ii) may be present.

We now analyze these three cases separately.

Case (i). This is similar to the analysis of Section IV. The optimal strategy for DM1, which leads to (u_1^t, u_2^t) as final outcome, independent of possible deviations in DM2's perception of α^0 , is given by

$$\gamma_1^{(i)}(u_2) = u_1^t + P^{(i)}(u_2 - u_2^t) \quad (29)$$

where

$$P^{(i)} = \frac{S_2 Q_2 [\beta_2 x_{20} + u_1^t - S_2 (\beta_1 x_{10} + u_1^t) - v_2]}{R_2 (u_2^t - v_2) + Q_2 (\beta_2 x_{20} + u_2^t - S_2 (\beta_1 x_{10} + u_1^t) - v_2)} \quad (30)$$

Case (ii). The solution here again follows from the analysis of Section IV. Hence, there exists an optimum insensitivity strategy realizing the team solution independent of DM2's different perceptions of Q_2 , and such a strategy is given by

$$\gamma_1^{(ii)}(u_2) = u_1^t + \frac{1}{S_2 \beta_1} (u_2 - u_2^t) \quad (31)$$

Case (iii). This case involves multiparameters where condition (20) is not satisfied. Hence, within the class of affine policies, there does not exist any element which makes the cost of DMI completely insensitive to discrepancies in DM2's perceptions in more than one parameter. In order to overcome this difficulty, we adopt, as discussed in Section V, the scalarized sensitivity function $\text{Tr}(I_2(\alpha^0, Q_2^0))$, and minimize it subject to the constraint that the strategy of DMI is given by

$$\gamma_1^{(iii)}(u_2) = u_1^t + P^{(iii)}(u_2 - u_2^t) \quad (32)$$

This problem can be shown to admit a unique solution which can be obtained explicitly. Hence, when DMI is uncertain about DM2's perception of both α and Q_2 , there still exists an affine strategy which minimizes an appropriate scalar function representing the sensitivity of DMI's incurred cost with respect to deviations in these coefficients from their nominal values, and such a strategy is given by (32).

In the preceding analysis, $P^{(i)}$ is the same coefficient as DMI would have used in his strategy in a Stackelberg game with DM2 being the follower and DMI enforcing the point (u_1^t, u_2^t) . On the other hand, in case (ii), by announcing a strategy of the form (31), DMI makes DM2's objective functional independent of the uncertain coefficient Q_2 . Therefore, DM2's discrepancies do not affect the team solution anymore. However, when the number of uncertain coefficients is large as compared with the dimension of DMs' decision vectors, there still exists a compromise, which is to minimize the cumulative effect of variations of uncertain parameters around their nominal values: $\gamma_1^{(iii)}(u_2)$ is designed to perform such a compromise.

VII. CONCLUDING REMARKS

In this paper we have introduced the notion of optimum minimum sensitivity incentive policies in team decision problems wherein one member of the team has a somewhat different perception of the common goal than the other one, and we have derived explicit incentive policies which render the incurred value of the team objective functional least sensitive to, and in some cases even independent of, the discrepancies described above.

One field where this notion finds application is the military Command, Control, and Communications (C³) systems area. Here, there exist multiple decision makers (DM's) and multiple hierarchies in decision making, and the role of each DM is not only to issue orders to be executed by the DM's occupying the lower levels of hierarchy, but also to coordinate the actions

of his subordinates. There is an underlying goal, or objective, which involves a successful completion of a mission or task (such as multi-object tracking and fire control), and this goal is determined by the DM's at the top of the hierarchy in rather general terms (i.e., not in fine detail), which is then transmitted to the relevant DM's at the lower levels.

Hence, in a general framework, a C³ system involves a team of DM's who act in an uncertain environment, and who have limitations on control and communication capabilities. However, realistically, this is not strictly a team problem, because, in an uncertain environment, it is unlikely that every DM will develop precisely the same perception of the ultimate goal in every fine detail. In fact, in order to model C³ systems as team problems, it is absolutely necessary that all DM's have exactly the same perception of an existing common goal and quantify this perception in exactly the same way. Any discrepancy that exists between the perceptions of the DM's on the underlying common goal will lead to a decision problem which cannot be treated as a team problem, and optimal decision rules derived by totally ignoring (or overlooking) this aspect of the problem are apt to lead to outcomes which are extremely sensitive even to small variations in the perceptions with regard to real underlying goals of the mission. The approach developed in this paper remedies this deficiency because it takes into account the possibility that the DMs' perceptions of the "team goal" may deviate from the nominal set by the highest level decision making unit.

Two possible extensions of the general approach of this paper are to dynamic multi-stage decision problems and to stochastic team problems. In the latter case a natural source of discrepancy is the a priori statistical information which is normally assumed to be shared by the DM's. A recent reference [Başar (1983)] addresses the question of existence of suitable equilibrium solutions for such problems when there is discrepancy in the subjective probability measures characterizing the probability space. Derivation of minimum sensitivity incentive policies in this context is currently under study.

REFERENCES

- Başar, T. (1983). Decision problems with multiple probabilistic models. *Proc. American Control Conf.*, San Francisco, CA, pp. 1091-1096.
- Cansever, D. H. and T. Başar (1982). A minimum sensitivity approach to incentive design problems. *Proc. 21st IEEE Conf. on Decision and Control*, Orlando, FL, pp. 158-163.
- Cansever, D. H., T. Başar, and J. B. Cruz, Jr. (1983). Robustness of incentive policies in team problems with discrepancies in goal perceptions. *Proc. 4th IFAC/IFORS Conf. on the Modelling and Control of National Economies*, Washington, D.C., pp. 275-280.
- Ho, Y. C., P. H. Luh, and G. J. Olsder (1982). A control theoretic view on incentives. *Automatica*, Vol. 18, No. 2, pp. 167-180.
- Marschak, J. and R. Radner (1972). *Economic Theory of Teams*, Yale University Press, New Haven, CT.
- Richardson, L. F. (1960). *Arms and Insecurity*. The Boxwood Press, Chicago, IL.

Salmas, M. and J. B. Cruz, Jr. (1981). An incentive model of duopoly with government coordination. Automatica, Vol. 17, pp. 821-830.

Schmitendorf, W. E. and G. Leitmann (1974). A simple derivation of necessary conditions for Pareto optimality. IEEE Trans. on Automatic Control, Vol. AC-19, pp. 601-602.

Simaan, M. and J. B. Cruz, Jr. (1975a). Formulation of Richardson's model of arms race from a differential game viewpoint. The Review of Economic Studies, Vol. XLII, pp. 67-77.

Simaan, M. and J. B. Cruz, Jr. (1975b). Nash equilibrium strategies for the problem of armament race and control. Management Science, Vol. 22, pp. 96-105.

Zheng, Y. P. and T. Başar (1982). Existence and derivation of optimal affine incentive schemes for Stackelberg games with partial information: A geometric approach. International Journal of Control, Vol. 35, No. 6, pp. 997-1011.

Peter B. Luh*

Taikang Ning*

Tsu-Shuan Chang**

Abstract A hierarchical decision problem with one leader and a continuum of followers is investigated. The problem is formulated as a stochastic Stackelberg game with nonnested information structure, and studied by using the Inducible Region concept. For the single-stage case, we show that the inducible region can be delineated, and the optimal Stackelberg strategy can be constructed. These results are then extended to the two-stage case. Although the problem is formulated in a pricing context in terms of companies and customers, the formulation can be interpreted differently to model C related problems.

1 Introduction

1.1 Motivation

The leader-follower type of problems, also known as the Stackelberg games, deal with multi-person hierarchical decision problems. Decision makers in such a problem belong to different levels of a hierarchy and have asymmetric roles. A higher level decision maker (a leader) has the authority to announce his strategy and impose it on the decision makers under him (the followers). The followers, knowing the leader's strategy, are assumed to react rationally to minimize their own cost functions. The leader's objective is then to design his strategy, by taking into account the followers' reactions, to minimize his cost. Since our society is essentially hierarchical in structure, there exist ample instances which can be modeled by such a framework. Among them are resource allocation, organization theory, pricing problems, and military command, control and communication (C³) systems.

As an example, consider the problem of platform positioning in a naval battle group. Ships in a naval battle group have various levels of defensive capability against enemy's air, surface and subsurface threats. The overall defensive effectiveness of the battle group depends not only on the capabilities of its member ships, but also on the positioning of all the ships (CAS82). In one of the doctrines, the ships are partitioned into subgroups for different defensive functions, such as anti-air and anti-submarine. Each subgroup has a function commander who is required to position his ships to maximize his functional defensive effectiveness. The overall defensive effectiveness is then determined by these functional measures. The allocation of the ships to subgroups by the top commander, the coordination and competition of the function commanders to maximize their own or the group performance, can be described as a hierarchical decision problem.

As mentioned, the leader's major task in such a problem is the design of his strategy, by taking

into account the follower's reactions, to minimize his cost. Mathematically, this implies the presence of composite functions in the leader's optimization problem (as will become clear in subsection 2.2). Any direct solution to such a problem seems almost impossible, except for special cases. Most of the recent advances are due to the discovery of two indirect approaches, the team solution approach (<BAS79,81>, <TOL81a>, <HO82>, <ZHE82>), and the Inducible Region approach (<CHA82a,82b,83>, <LWH83,84>, <TOL81b,83>).

In this paper, we use the inducible region concept to study a class of Stackelberg games, where there is one leader and a continuum of followers. Our purpose is not to model and solve a realistic, or even simplified C³ problem. Rather, we shall adopt a very basic model to study hierarchical decisionmaking, and demonstrate how it can be solved. For some reasons, the problem is formulated in a pricing context in terms of companies and customers. It should be noted, however, that the formulation can be interpreted differently to model C³ related problems.

1.2 The Pricing Problem and Outline of the Paper

Consider a pricing problem in which a company sells a certain kind of product to its customers. The company has to design its pricing scheme, and make it known to the customers. Knowing the pricing scheme, each customer then decides the amount to purchase. Therefore the problem can be viewed as a Stackelberg game, with the company as the leader and customers as followers. One way to model a large number of customers is by means of a histogram based on a key parameter such as customer's valuation of the product. After normalization, the histogram can be thought of as a probability density function. With the customer population being so described, we have a stochastic problem. If price differential is not allowed, the company can not observe and make use of customers' private information (individual's valuation of the product). We therefore end up with a non-nested stochastic Stackelberg Game. This model can also be interpreted as though there is a single follower who possesses some private information not known to the leader. Problems of this class have been treated to some extent by economists, e.g., <WEI74>, <SPE77>. There is no existing results in the control or game literature.

In Section 2, we study a single-stage pricing problem. The model is formulated in subsection 2.1. In subsection 2.2, we define the inducible region for such a stochastic game. We then in subsection 2.3 go through a sequence of steps to delineate the inducible region, and provide a systematic way to solve such a problem.

In Section 3, we study a two-stage pricing problem where a company sells a product to a continuum of customers during two consecutive time periods. The dynamic information involved, and the substitution/complement effects (to be defined in subsection 3.2) of the product at different time periods make the problem more complicated. The formulation is given in subsection 3.1. In subsection 3.2, we characterize the customers' reactions to a given pricing scheme.

* The research work reported here was supported in part by the National Science Foundation under Grant ECS 8105984, Grant ECS 8210673, and by the Air Force Office of Scientific Research under Grant 80-0098.

+ Department of Electrical Engineering and Computer Science, University of Connecticut, Storrs, CT 06268.

** Department of Electrical Engineering, SUNY at Stony Brook, Stony Brook, NY 11794.

The inducible region for the complementary case is then delineated in subsection 3.3, where we also show how to find optimal consumption curves, and optimal incentive pricing schemes. The substitution case is investigated in subsection 3.4. Concluding remarks and comparison to Spence's work are given in Section 4.

2 Pricing Problems in Non-nested Stochastic Stackelberg Games

2.1 Problem Formulation

Consider a pricing problem where there is a company (leader, DM0) that sells a product to a continuum of customers (followers). The customers are indexed by a parameter z , indicating individuals' valuation of goods. Assume that $z \in Z = [z_0, z_1]$, and is described by a probability density function $P_z(z)$ known to the leader. A customer with parameter z is denoted as DM1. Let r_0 denote the company's pricing scheme. If DM1 purchases u_1 (>0) units of the product, he has to pay $r_0(u_1)$ dollars. The social gain can be represented as

$$J_1(u_1, z) = S(u_1, z) - r_0(u_1). \quad (2.1.1)$$

The function $S(u_1, z)$ is DM1's satisfaction in dollar value by purchasing u_1 units of the product. We assume that S is twice differentiable with $S_{u_1 z} > 0$, so that a customer with higher valuation of goods has higher marginal satisfaction at every level of u_1 . The function $r_0(u_1)$ is assumed to be twice differentiable for $u_1 \geq 0$. For a given r_0 , DM1 decides an optimal u_1 to maximize J_1 , i.e.,

$$\max_{u_1} \{ S(u_1, z) - r_0(u_1) \} \quad \text{subject to } u_1 \geq 0. \quad (2.1.2)$$

The solution of (2.1.2) is denoted as $u_1^0(z)$. It represents DM1's "reaction" to the given r_0 . The price he has to pay for it is $u_1 = r_0(u_1^0(z))$. For simplicity, we assume that $u_1^0(z)$ exists and is unique. If $u_1^0(z)$ is not unique, DM0 has to assume the worst case in his strategy designing stage. This complicates the derivation, yet adds very little understanding to the problem.

The company's payoff function is described by

$$J_0 = \int_{z_0}^{z_1} L_0(r_0(u_1), u_1(z), z) P_z(z) dz. \\ = E \{ L_0 \}. \quad (2.1.3)$$

J_0 can be the profit, a social welfare function, or any meaningful function from the company's viewpoint. For the moment, however, we shall not specify J_0 explicitly. Knowing the density function $P_z(z)$ and followers' rationale in reaction ((2.1.2)), DM0 wants to select a strategy r_0 from some admissible set Γ_0 to maximize J_0 . In other words, DM0 solves the following problem:

$$(P-1) \max_{r_0 \in \Gamma_0} E \{ L_0(r_0(u_1), u_1(z), z) \} \quad \text{subject to } (2.1.2). \quad (2.1.4)$$

The decision sequence is summarized as follows.

$P_z(z)$	$r_0(\cdot)$	$u_0 = r_0(u_1)$
		$u_1^0(z)$
DM0's Prior Information	Design Stage	Execution Stage

Fig. 2.1

Problem (P-1) is a single-stage, stochastic

Stackelberg game. The information structure is non-nested, since each follower knows his own parameter z , but the leader can not observe and make use of it. The model, depending on the functional form of L_0 , can also be interpreted as though there is a single follower who possesses some private information z not known to the leader.

2.2 Inducible Region for Stochastic Games

Problem (P-1) is generally an intractable problem due to the presence of the composite function $r_0(u_1^0(z))$, and the non-nested information structure. In this section, we shall define the Inducible Region for this stochastic problem, and use it to solve the pricing problem.

For a given r_0 , each follower solves (2.1.2) to obtain $u_1^0(z)$. The mapping from Z to U_1 generated by $u_1^0(z)$ is denoted as r_1 . We then say that the r_1 is induced by this r_0 , and (r_0, r_1) is an inducible pair of strategies. For a different DM0's strategy, the followers' reactions will be changed accordingly, and we have another inducible pair of strategies. The Inducible Region in the Strategy Space, IRS, is then defined as the collection of all the inducible pairs of strategies, i.e.,

$$IRS = \{ (r_0, r_1) : (r_0, r_1) \text{ is an inducible pair of strategies} \}. \quad (2.2.1)$$

Since any realizable strategy pair must belong to IRS, (P-1) is equivalent to the following problem:

$$(P-2) \max_{(r_0, r_1) \in IRS} E \{ L_0(r_0, r_1, z) \}. \quad (2.2.2)$$

In (P-2), there is no composite function. However, IRS is a set of pairs of functions and is hard to be delineated. Furthermore, even if IRS can be delineated, it is still not easy to perform the maximization, since it involves r_0 and r_1 at the same time. Motivated by results of deterministic Stackelberg games, we shall next find IRS's counterpart in the decision space.

Consider a pair $(r_0, r_1) \in IRS$. DM1's decision is given by $u_1 = r_1(z)$. The corresponding u_0 is determined by $r_0(z) = r_0(r_1(z)) = f_0(z)$, where $f_0(z)$ is defined as the composite function of r_0 and r_1 that maps Z into U_0 . The pair $(u_0 = f_0(z), u_1 = r_1(z))$ can be regarded as the outcome for DM1 for this r_0 . For a follower with parameter z' , the corresponding outcome is given by $(f_0(z'), r_1(z'))$. Thus across the population, the outcomes are described by the pair of functions (f_0, r_1) . We then define IRD, inducible region in the decision space, as

$$IRD = \{ (f_0, r_1) : \exists (r_0, r_1) \in IRS \text{ s.t. } f_0(z) = r_0(r_1(z)) \quad \forall z \in Z \}. \quad (2.2.3)$$

It is clear that (P-2) is equivalent to the following optimization problem:

$$(P-3) \max_{(f_0, r_1) \in IRD} E \{ L_0(f_0, r_1, z) \}. \quad (2.2.4)$$

where appropriate substitutions of r_0 by f_0 in L_0 are assumed. For deterministic problems, IRD is a subset in the decision space ($\langle CP^4 82b \rangle, \langle LM 83, 84 \rangle$). For our problem, however, due to the presence of a continuum of customers, IRD remains to be a set of pairs of functions.

2.3 The Inducible Region Approach 2.3.1 Customers' First and Second Order Necessary Conditions

In order to delineate IRD and solve the problem, we shall first examine a customer's reaction to a pricing

eme. For a given r_0 , a customer has to decide whether to buy the product or not. Assume $r_0(0) = 0$, i.e., a customer pays nothing if he does not buy it. If

$$(0, z) > \{S(u_1, z) - r_0(u_1)\} \quad \forall u_1 > 0, \quad (2.3.1.1)$$

customer will not purchase any of the product, i.e., $u_1 = 0$. Let z_d be the value of z such that $u_1 = 0$ for $z \leq z_d$. Thus, z_d can be thought of as the "try" point.

Consider a customer with $z > z_d$. The first order necessary condition of (2.1.2) is

$$dj_1/du_1 = 0, \text{ or} \quad S_{u_1} = dr_0/du_1. \quad (2.3.1.2)$$

A solution will lead to the function $r_1(z)$. The second order necessary condition is

$$d^2j_1/du_1^2 = (S_{u_1u_1} - d^2r_0/du_1^2) \leq 0. \quad (2.3.1.3)$$

To get more explicit result, we take total derivative on both sides of (2.3.1.2) with respect to z . We have

$$(S_{u_1u_1}) dr_1/dz + S_{u_1z} = (d^2r_0/du_1^2) dr_1/dz, \text{ or} \\ \{S_{u_1u_1} - (d^2r_0/du_1^2)\} dr_1/dz = -S_{u_1z}. \quad (2.3.1.4)$$

Since $S_{u_1z} > 0$ by assumption, the right hand side of (2.3.1.4) is not zero. Therefore, neither $(S_{u_1u_1} - d^2r_0/du_1^2)$ nor dr_1/dz is zero. Equation (2.3.1.3) then implies that

$$d^2j_1/du_1^2 = \{-S_{u_1z}/(dr_1/dz)\} < 0, \text{ and} \quad (2.3.1.5)$$

$$dr_1/dz > 0. \quad (2.3.1.6)$$

thus have

Lemma 2.1: The customers' reactions $r_1(z)$ is a strictly increasing function of z for $z > z_d$ (i.e., a customer with higher valuation buys more goods). So, (2.3.1.2) and (2.3.1.5) constitute a set of sufficient conditions for customers with $z > z_d$.

Since $r_1(z)$ is strictly increasing for $z > z_d$, its inverse mapping exists, denoted as $z = r_1^{-1}(u_1)$. Thus, though the company can not observe directly a customer's private information z , however, by knowing his decision u_1 , the company can indirectly figure it out. This point, as a result of the assumption that $u_1 z > 0$, turns out to be crucial in our design of the company's pricing scheme.

3.2 Construction of the Composite Function f_0

Define the Set of Reactions, SR , as the collection of all inducible reactions, i.e.,

$$SR = \{r_1: \exists r_0 \text{ s.t. } (r_0, r_1) \in IRS\}, \text{ or} \\ = \{r_1: \exists f_0 \text{ s.t. } (f_0, r_1) \in IRD\}. \quad (2.3.2.1)$$

we have the following result.

Theorem 2.1: For any $r_1 \in SR$, its corresponding f_0 can be constructed within the limit of a constant term, and is given by

$$f_0(z) = S(r_1(z), z) - \int_{z_d}^z (\partial S / \partial z) dz' - C_1, \quad (2.3.2.2)$$

$$= \left(\int_{0}^{u_1} (r_1^{-1}(u_1'), u_1') du_1' + C_2 \right) \Big|_{u_1 = r_1(z)}. \quad (2.3.2.3)$$

where C_1 and C_2 are constants to be determined. Furthermore, one r_0 that induces the r_1 is

$$r_0(u_1) = \int_{0}^{u_1} (r_1^{-1}(u_1'), u_1') du_1' + C_2. \quad (2.3.2.4)$$

Due to limit on page size, proofs for Theorem 2.1 and remaining theorems are not included here. For their proofs, please see <NIN83>.

2.3.3 Solution Method

Theorem 2.1 is very useful since it says that for any $r_1 \in SR$, its corresponding f_0 , as well as r_0 , can be constructed within the limit of a constant. Thus the optimization of J_0 in (P-3) with respect to (f_0, r_1) over the set IRD can essentially be reduced to the optimization of J_0 with respect to r_1 over the set SR . The constant term in f_0 is related to z_d , and can generally be resolved separately. We thus have a reduced version of (P-3) as follows:

$$(P-4) \quad \max_{r_1 \in SR} E\{L_0(r_1, z)\}, \quad (2.3.3.1)$$

where appropriate substitutions of f_0 by (2.3.2.7) or (2.3.2.9) are assumed. Note that (P-4) is a maximization of J_0 with respect to a single function r_1 over the set SR , rather than a maximization with respect to two functions as in (P-3). The next theorem then delineates SR .

Theorem 2.2: The Set of Reaction SR is given by

$$SR = \{r_1: dr_1/dz > 0 \text{ for } z > z_d, \text{ and } r_1(z) = 0 \\ \text{for } z \leq z_d, z_d \in [z_m, z_M]\}. \quad (2.3.3.2)$$

Problem (P-4) can be rewritten as follows:

$$(P-4') \quad \max_{r_1} E\{L_0(r_1, z)\} \text{ s.t. } dr_1/dz > 0 \text{ for } z > z_d, \text{ and} \\ r_1(z) = 0 \text{ for } z \leq z_d, z_d \in [z_m, z_M]. \quad (2.3.3.3)$$

It is an "optimal control like" problem with an inequality constraint. Although its solution is nontrivial and subject to further investigation, nevertheless, we have converted the original intractable problem (P-1) to the comparatively much simpler problem (P-4), where its solution, or at least a solution, is promising. Once (P-4) is solved, its result r_1^* is then the best inducible r_1 . The optimal pricing strategy can then be constructed according to (2.3.2.8). For completeness, the entire inducible region IRD is presented as follows.

Theorem 2.3: IRD is delineated by

$$IRD = \{(f_0, r_1): r_1 \in SR, \text{ and } f_0 \text{ is given by (2.3.2.7) or (2.3.2.9)}\}.$$

3 Two-Stage Pricing Problems

3.1 Problem Formulation


In this section, we extend the model and the Inducible-Region approach of Section 2 to study a two-stage pricing problem, where a company sells a product to a continuum of customers at two consecutive time periods. In many situations, once a customer decides to enter a market, he has to pay an entrance fee regardless his consumption level. For instance, a telephone company charges a basic service fee to a customer in the monthly bill even though he did not

use the telephone for the entire month. We shall therefore assume

$$r_0 = \begin{cases} r_{01} + r_{02} + A & \text{for committed customer,} \\ 0 & \text{otherwise,} \end{cases} \quad (3.1.1)$$

where r_{01} denote the company's pricing strategy at stage 1, and A the entrance fee. Let $u_{11}(z)$ denote the amount of the product purchased by DM_1 at stage 1 ($i=1,2$). Also, let y_i denote the on-line uncertainty at stage 1 such as weather that affect customers' satisfaction as well as the company's payoff. It is described by the probability density functions $P_{y_i}(y_i)$. For simplicity, we assume that y_1 , y_2 and z are mutually independent. Finally, similar to Section 2, we let r_{11} denote the consumption curve at stage 1. The sequence of actions is summarized in Fig. 3.1. In this figure, g_{11} and g_{21} denote, respectively, the information available to the company and DM_1 at stage 1, and are given by

$$\begin{aligned} g_{21} &= (z, y_1; r_{01}, r_{02}), \\ g_{p1} &= (y_1, u_{11}), \\ g_{22} &= (z, y_1, u_{11}, y_2; r_{01}, r_{02}), \\ g_{p2} &= (y_1, y_2, u_{12}). \end{aligned} \quad (3.1.2)$$



We let $S(z, y_1, y_2, u_{11}, u_{12})$ represent DM_1 's satisfaction in dollar value, and the social gain z is given by

$$J_1 = S(z, y_1, y_2, u_{11}, u_{12}) - r_{01}(y_1, u_{11}) - r_{02}(y_1, y_2, u_{12}) - A. \quad (3.1.3)$$

We assume that r_{01} and S are twice differentiable for $u_{11} > 0$, with $S_{u_{11}} > 0$ and $S_{u_{12}} > 0$, i.e., a customer with larger z has higher marginal satisfaction for every level of purchase at both stages. We also assume for simplicity that r_{01} is continuous at $u_{11} = 0$ with value zero if $u_{11} = 0$.

For a given pricing scheme (r_{01} , r_{02} and A), customers decide their optimal consumption strategies by maximizing J_1 , i.e.,

$$\max_{r_{11}, r_{12}} E \{ J_1 \}. \quad (3.1.4)$$

Let J_1^0 be the payoff of DM_1 if he does not enter the market, i.e.,

$$J_1^0 = E \{ S(z, y_1, y_2, u_{11}=0, u_{12}=0) \}. \quad (3.1.5)$$

Also let J_1^p be the payoff of DM_1 if he enters the market, i.e.,

$$\begin{aligned} J_1^p &= E \{ S(z, y_1, y_2, u_{11}, u_{12}) - r_{01}(y_1, u_{11}) - r_{02}(y_1, y_2, u_{12}) - A \}, \\ &\text{with } u_{11} \geq 0 \text{ and } u_{12} \geq 0. \end{aligned} \quad (3.1.6)$$

DM_1 will enter the market if $J_1^p > J_1^0$ for some $u_{11} \geq 0$ and $u_{12} \geq 0$. Let z_d denote the "entry point", i.e., a customer with $z > z_d$ will enter the market. For simplicity of discussion, we shall assume in this section that $z_d > z_c$. Then since $J_1^p = J_1^0$ at $z = z_d$, A and z_d are related by the following equation:

$$\begin{aligned} A &= E \{ S(z_d, y_1, y_2, u_{11}, u_{12}) - r_{01}(y_1, u_{11}) - r_{02}(y_1, y_2, u_{12}) \} \\ &\quad - E \{ S(z_d, y_1, y_2, 0, 0) \}. \end{aligned} \quad (3.1.7)$$

Knowing the customers' rationale ((3.1.4)), the company then designs r_{01} , r_{02} and A to maximize his payoff J_0 , where

$$J_0 = E \{ L_0(r_{11}, r_{12}, r_{01}, r_{02}) \}. \quad (3.1.8)$$

That is, the company solves the following optimization problem:

$$(P-1) \quad \max_{r_{01}, r_{02}, A} \{ J_0 \} \quad \text{subject to (3.1.4).}$$

The problem is thus a two-stage, stochastic Stackelberg game. Similar to the single-stage pricing problem of Section 2, the information structure is non-nested since the company can not observe and make use of the information z . We shall next solve this problem by extending the inducible region approach of Section 2.

3.2 Reactions of Customers

For simplicity, we shall assume that once a customer decides to enter the market, he will purchase the product at both periods (i.e., $u_{11} > 0$ and $u_{12} > 0$). These customers solve a two-stage dynamic optimization problem, i.e.,

$$\max_{r_{12}} E \{ J_1(r_{11}, r_{12}, r_{01}, r_{02}) | g_{22} \}, \quad \text{and} \quad (3.2.1)$$

$$\max_{r_{11}} E \{ J_1(r_{11}, r_{12}^*, r_{01}, r_{02}) | g_{21} \}, \quad (3.2.2)$$

where r_{12}^* denotes the solution of (3.2.1). For simplicity, we assume that both (3.2.1) and (3.2.2) have unique solutions. To better characterize customers' reactions for a given pricing scheme, we shall first examine necessary conditions of (3.2.1) and (3.2.2). The first and second order necessary conditions of (3.2.1) are as follows:

$$dJ_1/du_{12} = 0, \quad \text{and}$$

$$d^2J_1/du_{12}^2 \leq 0 \quad \text{for } u_{12} > 0.$$

Or,

$$S_{u_{12}} = \partial r_{02} / \partial u_{12}, \quad \text{and} \quad (3.2.3)$$

$$(S_{u_{12}u_{12}} - \partial^2 r_{02} / \partial u_{12}^2) \leq 0 \quad \text{for } u_{12} > 0. \quad (3.2.4)$$

Note that there is no conditional expectation in the above equations since at the second period, g_{22} includes everything. Similarly, the necessary conditions for (3.2.2) are

$$d(E_{y_2} \{ J_1 \}) / du_{11} = 0, \quad \text{and} \quad (3.2.5)$$

$$d^2(E_{y_2} \{ J_1 \}) / du_{11}^2 \leq 0 \quad \text{for } u_{11} > 0. \quad (3.2.6)$$

The expectations are taken over y_2 , i.e., the on-line uncertainty at stage 2. From (3.2.5) and (3.2.3), we have

$$\begin{aligned} E_{y_2} \{ S_{u_{11}} - \partial r_{01} / \partial u_{11} + (S_{u_{12}} - \partial r_{02} / \partial u_{12}) dr_{12} / du_{11} \} \\ = E_{y_2} \{ S_{u_{11}} - \partial r_{01} / \partial u_{11} \} = 0. \end{aligned} \quad (3.2.7)$$

Therefore, (3.2.5) becomes

$$E_{y_2} \{ S_{u_{11}} \} = \partial r_{01} / \partial u_{11}. \quad (3.2.8)$$

Similarly, we can expand (3.2.6) to obtain

$$E_{y_2} \{ S_{u_{11}u_{11}} + S_{u_{11}u_{12}} (\partial r_{12} / \partial u_{11}) - \partial^2 r_{01} / \partial u_{11}^2 \} \leq 0. \quad (3.2.9)$$

We have

Lemma 3.1: If $(S_{zu12} + S_{u1u12}(dr_{11}/dz)) > 0$, then $dr_{11}/dz > 0$ for $z > z_d$, i.e., $r_{11}(z)$ is a strictly increasing function of z .

Similarly, we have

Lemma 3.2: If $E_{z2}(S_{zu11} + S_{u1u12}(dr_{12}/dz)) > 0$, then $dr_{12}/dz > 0$ for $z > z_d$, i.e., $r_{12}(z)$ is a strictly increasing function of z .

Lemma 3.1 and Lemma 3.2 are useful since if they hold, then both r_{11} and r_{12} are strictly increasing functions of z , and their inverse mappings $r_{11}^{-1}(u_{11}; y_1)$ and $r_{12}^{-1}(u_{12}; y_2)$ exist. Frequently, knowing a customer's decision u_{11} or u_{12} , the company can figure out his parameter z . Similar to the single-stage case, this point is useful in designing the company's pricing schemes. Fully utilize Lemma 3.1 and Lemma 3.2, we have to examine the conditions under which they are satisfied. We shall first characterize the term dr_{12}/dz in Lemma 3.2.

Lemma 3.3: If $S_{zu12} > 0$, then $dr_{12}/dz > 0$ for $z > z_d$.

Note that the term S_{u1u12} appears in both Lemma 3.1 and Lemma 3.2. We shall next consider the case where

$$S_{u1u12} > 0, \quad (3.2.19)$$

i.e., a customer's marginal satisfaction at one stage is enhanced as the purchased amount at the other stage is increased. In other words, the product at different time periods have complement effect. The other situation when the substitution effect exists, i.e., $S_{u1u12} < 0$, will be discussed in Section 3.4. It can be easily checked that if $S_{u1u12} > 0$, then since $S_{zu11} > 0$ and $S_{zu12} > 0$ by assumption, the conditions of Lemma 3.1 and Lemma 3.2 hold. We therefore have the following result.

Theorem 3.1: For the two-stage pricing problem as formulated, if $S_{u1u12} > 0$, then r_{11} and r_{12} are strictly increasing functions of z for $z > z_d$.

3.3 Inducible Region and Stackelberg Strategies

In this subsection, we shall first show by constructive proof that for any pair (r_{11}, r_{12}) of strictly increasing functions of z , there exists a pair of incentive pricing schemes to induce it. We then delineate the entire inducible region for the case when $S_{u1u12} > 0$.

Theorem 3.2: For any pair of functions (r_{11}, r_{12}) with $dr_{11}/dz > 0$ and $dr_{12}/dz > 0$ for $z > z_d$, there exist a pair of incentive pricing schemes $(r_{01}, r_{02}) \in \Gamma_{01} \times \Gamma_{02}$ to induce them. They are given by

$$r_{01} = \int_{u_{11}}^{u_{11}'} (r_{11}^{-1}(y_1, y_2, u_{11}', r_{12}(r_{11}^{-1}(y_1, y_2, u_{11}')), u_{11}')) du_{11}' \quad (3.3.1)$$

and

$$r_{02} = \int_{u_{12}}^{u_{12}'} (r_{12}^{-1}(y_1, y_2, r_{11}(r_{12}^{-1}(y_1, y_2, u_{12}')), u_{12}')) du_{12}' \quad (3.3.2)$$

The required entrance fee A is given by (3.1.7).

Similar to the single-stage case, we can define IRS, the Inducible Region in the Strategy Space, and SR, the Set of Reactions. We then have the following

result:

Theorem 3.3: For the case where $S_{u1u12} > 0$, the Set of Reactions SR is given by

$$SR = \{(r_{11}, r_{12}) : dr_{11}/dz > 0 \text{ and } dr_{12}/dz > 0 \text{ for } z > z_d, \text{ and } r_{11}(z) = r_{12}(z) = 0 \text{ for } z \leq z_d, z \in [z_m, z_H]\}. \quad (3.3.3)$$

Now similar to Section 2, the original intractable problem (P-1) is equivalent to

$$(P-2) \max_{SR} E\{L_0(r_{11}, r_{12})\}. \quad (3.3.4)$$

where appropriate substitutions of r_{01} and r_{02} by (3.3.1) and (3.3.2) are assumed. Furthermore, z_d , rather than A , is treated as an independent parameter to be optimized. Consequently, the remaining problem can be solved by going through the following steps:

- (1) find necessary conditions for r_{12} and r_{11} by solving an optimal control like problem at each period,
- (2) find necessary conditions for z ,
- (3) solve conditions derived from (1) and (2) simultaneously or sequentially, and
- (4) construct the optimal incentive pricing schemes according to (3.3.1), (3.3.2) and (3.1.7).

Finally, for completeness, we have

Theorem 3.4: For goods with complement effect, IRS is delineated by

$$IRS = \{(A, r_{01}, r_{02}, r_{11}, r_{12}) : (r_{11}, r_{12}) \in SR, A, r_{01}, r_{02} \text{ are given by (3.1.7), (3.3.1) and (3.3.2), respectively}\}. \quad (3.3.5)$$

3.4 The Case of Substitutive Goods

In this section, we shall discuss the case where the product at different time periods have substitution effect. It can be checked that now the conditions of Lemma 3.1 and 3.2 might not be satisfied. Therefore, the results of the previous section do not hold in general. We shall impose appropriate conditions on both the satisfaction function S and the pricing strategies (r_{01}, r_{02}) , so that those conditions remain valid.

Consider an S with

$$S_{zu11} = q_1 > 0, \quad S_{zu12} = q_2 > 0, \quad (3.4.1)$$

$$S_{u1u12} = -q_3 < 0, \quad (3.4.2)$$

$$S_{u1u11} = -q_4 < 0, \text{ and } S_{u1u12} = -q_5 < 0. \quad (3.4.3)$$

The negativity of S_{u1u12} implies that the product at different time periods are substitutes. Equation (3.4.3) says that, at each stage, a customer's marginal satisfaction decreases as the amount of purchase increases. We then have

Theorem 3.5: For an S satisfying (3.4.1) to (3.4.3), suppose that

$$q_5 > (q_2/q_1) q_3, \text{ and} \quad (3.4.4)$$

$$q_4 > (q_1/q_2) q_3. \quad (3.4.5)$$

Then with affine pricing schemes $(r_{01}$ being affine in u_{11}), the customers' reactions are strictly increasing functions of z .

Intuitively, (3.4.4) and (3.4.5) imply that the substitution effect is weak. To see this, we first assume that $q_1 \geq q_2$. In this case, (3.3.5) and (3.4.5) become $q_3 < q_5$ and $q_3 < q_4$. That is, the substitution effect (q_5) is weaker than the decreasing in marginal satisfaction within each period (recall that $-q_4 = \partial S / \partial u_{11}$, and $-q_5 = \partial S / \partial u_{12}$). If $q_1 < q_2$, (3.3.5) and (3.4.5) then imply that q_3 should be weaker than some scaled version of q_4 and q_5 .

Under the conditions of Theorem 3.5, r_{11} and r_{12} are strictly increasing functions of z . Although (3.3.1) and (3.3.2) can be applied to obtain certain pricing functions, however, these functions might not be affine, thus the original assumption might be violated, and there might not exist a pair of affine pricing schemes that induces them. Therefore the result we have here is not as strong as those in the complementary case.

4 Concluding Remarks

Pricing problems are formulated as stochastic Stackelberg games with non-nested information structure, and studied by using the Inducible Region concept. It was thought that such problems were extremely difficult to solve, and their inducible regions are not clearly defined. However, by exploiting the special structure of the model under consideration, we are able to delineate the complete inducible region for the single-stage pricing problem. The optimal inducible consumption curve can then be obtained by solving an optimal control like problem. The corresponding optimal pricing scheme can also be constructed.

Spence studied single-stage pricing problems in [SPE77]. There is certain mathematical resemblance between his results and the results presented in Section 2. Some of our ideas are indeed motivated by his work. In [SPE77], he first substitutes out all the r_0 's by (2.3.2.2), and then maximizes J_0 over the unrestricted set of all the r_1 's. On the other hand, in the inducible region approach, we first substitute out r_0 's by (2.3.2.2) or (2.3.2.3) (they are shown to be equal). Instead of directly maximizing J_0 over the unrestricted set of all the r_1 's, we delineate the set of all the customers' reactions (SR). We then find the optimal inducible reaction by maximizing J_0 over SR. Once r_1^s is found, r_0^s can then be constructed from (2.3.2.4). Thus using the inducible region approach, we guarantee the solution obtained from (P-4) is inducible, however, this is not generally true by using Spence's method. For problems treated in [SPE77], his solutions happened to be in SR, thus he was able to get the results. For general problems, this will not be the case. We then have to rely on the inducible region concept.

For the two-stage pricing problem of Section 3, the dynamic information involved makes it more complicated than the single-stage problem, and the complement/substitution effect of the product at two different periods begins to emerge. If the product at different periods have complement effect, we delineate the entire Inducible Region. The desired consumption curves can be found by solving an optimal control like problem at each period, and the optimal pricing strategies can then be constructed. On the other hand, if the product has "weak" substitution effect, we show that the optimal affine pricing schemes can sometimes be constructed.

Appendix

[BAS79] Tamer Basar, Hassan Solbuz, "Closed-Loop Stackelberg Strategies with Application in the Optimal Control of Multilevel Systems," IEEE Transactions on Automatic Control, Vol. AC-24, No. 2, April 1979, pp. 166-179.

[BAS81] Tamer Basar, "Equilibrium Strategies in Dynamic Games with Multi-Level of Hierarchy," Automatica, Vol. 17, No. 5, 1981, pp. 749-754.

[CAS82] F.A. Castanon, J.R. Delaney, L.C. Kramer, M. Athans, "A Mathematical Framework for the Study of Battle Group Position Decisions," Proceeding of the 5th MIT/ONR Workshop on C² systems, Monterey, CA, Aug. 1982, pp. 105-110.

[CHA82a] Tsu-Shuan Chang, Peter B. Luh, "The Concept of Inducible Region in Stackelberg Games," the proceedings of the 1982 American Control Conference, Arlington, Virginia, June 1982, pp. 139-140.

[CHA82b] T. S. Chang, P. B. Luh, "A Complete Solution for Two-Person, Single-Stage, Deterministic, Stackelberg Games," Proceedings of the 21st IEEE Conference on Decision and Control, Orlando, Florida, December 1982, pp. 176-180.

[CHA83] Tsu-Shuan Chang, Peter B. Luh, "Derivation of Necessary and Sufficient Conditions for Single-Stage Stackelberg Game via the Inducible Region Concept," to appear in IEEE Transactions on Automatic Control, Nov. 1983.

[CHO82] Yi-Chi Ho, Peter B. Luh and G. J. Olsder, "A Control Theoretic View on Incentives," Automatica, Vol. 18, No. 2, March 1982, pp. 167-179.

[LUH83] Peter B. Luh, S. C. Chang, T. S. Chang, "Solutions and Properties of Multi-Stage Stackelberg Games," to appear in Automatica.

[LUH84] P. B. Luh, T. S. Chang, Taikang Ning, "Three-level Hierarchical Decision Problems," To appear in IEEE Transactions on Automatic Control, April 1984.

[NIN83] Taikang Ning, Master Thesis, Department of Electrical Engineering and Computer Science, The University of Connecticut, 1983.

[SPE77] M. Spence, "Nonlinear Prices and Welfare," Journal of Public Economics, Vol. 8, 1977, pp. 1-18.

[TOL81a] B. Tolwinski, "Closed-Loop Stackelberg Solution to Multi-Stage Linear-Quadratic Game," Journal of Optimization Theory and Application, Vol. 34, No. 4, August 1981, pp. 485-501.

[TOL81b] B. Tolwinski, "Equilibrium Solutions for a Class of Hierarchical Games," in Applications of Systems Theory to Economics, Management and Technology, J. Gitenbaum and M. Niezgodka, Eds., Warsaw, PWN, 1981, pp. 581-600.

[TOL83] B. Tolwinski, "A Stackelberg Solution of Dynamic Games," IEEE Transactions on Automatic Control, Vol. 28, No. 1, Jan. 1983, pp. 85-93.

[WEI74] M. L. Weitzman, "Prices vs. Quantities," The Review of Economic Studies, Vol. 41, No. 128, Oct. 1974, pp. 477-491.

[ZHE82] Y. P. Zheng, T. Basar, "Existence and Derivation of Optimal Affine Incentive Schemes for Stackelberg Games with Partial Information: A Geometric Approach," International J. of Control, Vol. 35, No. 6, 1982, pp. 997-1011.

OPTIMAL FILE ALLOCATION PROBLEMS FOR DISTRIBUTED DATA BASES IN UNRELIABLE COMPUTER NETWORKS II

by
Moses Hsingwen Ma and Michael Athans

Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

ABSTRACT

The problem to be investigated consists of determining the optimal locations of files and the number of redundant copies of these files, in a vulnerable ~~to~~ communications network. It is assumed that each node and link of the communication network can fail independently of the others. The optimization problem maximizes the probability that a commander is able to access the subset of the files that he needs while minimizing the network-wide costs. These network-wide costs are storage costs and costs due to the time delay in query and update requests of the distributed data base system.

The problem can be shown to reduce to a zero-one linear programming problem. We will look for theorems which reduce the complexity of the solution of the zero-one linear programs. Finally a heuristic algorithm has been developed to solve the zero-one linear program. An efficient polynomial time algorithm has been developed for the totally reliable network case. We will try to extend the efficient polynomial time algorithm to the unreliable network case.

destination and the assumption that each time a message enters a node a new length is chosen from the exponential message length distribution. The use of fixed routing is wrong since a link or node may not be operational at the time it is being used.

The Mahmoud and Riordon problem reduces to a nonlinear integer programming problem. They used a heuristic solution technique to solve the problem. The solution algorithm consisted of first generating a number of feasible initial solutions; second optimize this solution by successive additions or deletions of file copies; third adopt the lower cost allocations and repeat step two. Finally when additions or deletions cannot further reduce the cost, then this allocation is considered to be a "local optimum." Each time the algorithm produces a local optimum it is considered as an "heuristic run." The result of the lowest cost heuristic run is adopted as the optimal file allocation. The heuristic approach is assuming that one or more of the local optimal will be close in cost to the actual global optimum. The solution must satisfy the availability and time delay constraints.

1. INTRODUCTION

1.1. Literature Survey of Work Done on Unreliable Networks

Unreliable networks are networks in which the nodes or links may not be operational at any time due to failures. The first researchers to investigate the file allocation problem in an unreliable networks were Mahmoud and Riordon [1.1]. Mahmoud and Riordon developed a model in which not only the allocation of files were optimized, but also the optimal capacities of communication channels were determined. The interesting extension with respect

to previous models was the definition of availability constraints measuring the probability of having at least one copy of the file accessible as a function of link reliabilities, node reliabilities and network topology. Mahmoud and Riordon also used network delay constraints. They developed a model in which not only the allocation of files were optimized but also the optimal capacities of communication channels were determined. The delay measured used is from Kleinrock's [1.2] book "Stochastic Message flow and Delay". The delay expression assumes fixed routing, each node has infinite storage, intermediate nodes incur no processing delay as a message travels through it toward its

Finally file allocation in a distributed computer communication network with adaptive routing has been recently published by Lanning and Leonard [1.3]. They have presented an algorithm which minimizes storage and message transmission costs while the file allocations satisfy minimum file availability and message delay requirements. The algorithm uses the solution to a p-median problem to find an initial candidate file placement. The p-median problem is to allocate files at p locations so as to minimize all the average costs. If for example the costs were replaced by distance then the simple median problem would be to place the file in such a way that the average distance from a node to the file was minimized. For a p-median problem one would try to place the files in such a way that the average distance from a node to a copy of the file was minimized.

Lanning and Leonard used the p-median problem because he requires that the users place an upperbound on the number of file copies say p. Lanning's algorithm first solves the p median problem for each p up to the upperbound P. Second the algorithm examines each of the P solutions to the p-median problems and determines if the locations satisfy the availability constraints. If there are any infeasibilities the algorithm employs a set of rules to eliminate infeasibilities. Third the algorithm

examines the P solutions with respect to delay constraints. If there are any infeasibilities the algorithm again employs a set of rules to eliminate the infeasibilities. Once all the infeasibilities have been resolved the best of the P solutions is taken as the optimal solution. Lanning uses the same approximate availability measure as Mahmoud and Riordon for the availability constraints and uses a time delay approximation for the time delay constraints.

2. TIME DELAYS IN AN UNRELIABLE NETWORK

2.1. Basic Problem

The problem at hand is a completely unreliable system with significant time delays. Since the system is unreliable we will have unaccessability costs and because the system has time delays we will also have communication costs. Therefore we will have storage, communication and non-accessability costs in our objective function. We want to minimize this objective function subject to the fixed capacity constraints. Therefore our formulation for the problem at hand may be written in the following form:

$$\min_{I_l} \sum_{l=1}^M C(I_l) =$$

$$\min_{I_l} \sum_{l=1}^M \left[\sum_{j=1}^N \left[\sum_{k \in I_l} \psi_{jl} \bar{\epsilon}_{jk} + \min_{k \in I_l} \lambda_{jl} \zeta_{jk} \right] + \sum_{k \in I_l} \sigma_{kl} - \sum_{i=1}^L \alpha_i \beta_{il} P_{il}(I_l) - \sum_{i=1}^S \sum_{j=1}^S U_{ij} P_{ijl}(I_l) \right] \quad (2.1)$$

s.t.

$$x \geq 0$$

α_i = the importance of commander i
 β_{jl} = the value of file j to commander i
 $P_{il}(I_l)$ = the probability that file l is accessible to the commander at node i given an assignment I_l
 λ_{jl} = the volume of query traffic emanating from node j for file l
 ψ_{jl} = the volume of update traffic emanating from node j for file l
 $\bar{\epsilon}_{jk}$ = the time delay for an update from node j to node k
 ζ_{jk} = the time delay for a query from node j to node k
 σ_{kl} = the cost of locating a copy of a file j at the k^{th} node
 S = the number of sensors
 U_{ij} = the importance of file i when data from sensors i and j are fused
 $P_{ijl}(I_l)$ = the probability that file l is accessible sensors i and j given an assignment I_l
 I_l = the set of node indexes representing a given assignment of file l

2.2. Queueing Discipline

We require the most efficient queueing discipline to reduce time delays. In order to achieve this goal we must allocate higher priorities to the requests with shorter service times and lower priorities to requests with longer service times. In this manner the average waiting time per request is reduced. This is true because the request with longer service time must always wait for those requests with shorter service time, while the requests with shorter service times do not wait for requests with longer service times unless they are already being serviced. In this way requests

with shorter service times do not have to wait as long as a first come first serve (FCFS) queue. The requests with longer service times have a slightly longer waiting time than FCFS, but it is more than counterbalanced by the reduced waiting times for the requests with shorter service times. Therefore in our problem formulation queries have higher priorities than do updates since queries have shorter service times.

We also may have priorities due to other factors. These factors include importance of the commander, the importance of the file to the commander, the importance of the file to the multiple sensors. The waiting times due to the priorities can be calculated using formulas given by Kleinrock's [2.1] book entitled "Queueing Systems II".

2.3. Fixed Routing

In fixed routing the routing strategy remains fixed once it is determined. The set of routing variables, $\Phi_{ij}(k)$ define the routing that is to be done. $\Phi_{ij}(k)$ is the fraction of traffic emanating out of node i that is destined for node k on link (i,j). these set of routing variables can be found by some shortest path algorithm subject to link capacities.

Let δ_i be the generation rates of messages at node i, λ_i be the total amount of query traffic emanating out of node i, and ψ_i be the total amount of update traffic emanating out of node i. We then have the following:

$$\delta_i = \lambda_i + \psi_i \quad (2.2)$$

$$\sum_k \sum_j \Phi_{ij}(k) = 1 \quad (2.3)$$

$$\lambda_i = \sum_{l=1}^m \lambda_{il} \quad (2.4)$$

$$\psi_i = \sum_{l=1}^m \psi_{il} \quad (2.5)$$

Let $1/\mu$ be the average message length. Let θ_i be the total arrival rate of messages at node i . In other words θ_i is the sum of messages generated at node i and those in transit through node i . then we have the following relation:

$$\theta_i = \delta_i + \sum_j \sum_k \theta_j \Phi_{jk}(k) \quad (2.6)$$

The total arrival rate of messages to link (i,j) , denoted by Θ_{ij} , is given by:

$$\Theta_{ij} = \theta_i \sum_k \Phi_{ij}(k) \quad (2.7)$$

2.4. Variable Routing

The routing for this part of the formulation must be changed. In the previous sections we have used fixed routing strategies. In this section we must use variable routing strategies. The routing must be done in such a way to take into account failing links or changes in our topology of our network. Therefore due to link failures the routing variables are scaled up by the fraction of time they are operational. This is simple averaging of the routing variables with respect to the link failure probabilities.

The problem is now how to change this routing to effect the reliability of the network. Clearly the fixed routing strategy found in the previous paragraph will not do for our unreliable network. Therefore somehow these routing variables must be changed to reflect the reliability. We can use the method described by Li [2.2]. Li assumed that given a routing strategy that if a link failed that the remaining messages will flow through the other links proportionally. In this case we can then find new routing variables based on the old routing variables and the probability of link failures.

$$\Phi'_{ij}(k) =$$

$$\frac{\Phi_{ij}(k)\pi_{ij}}{\sum_{(i,j) \in E} \Phi_{ij}(k)\pi_{ij}} \quad (2.8)$$

where π_{ij} is the probability that link (i,j) is reliable and E is the set of all links in the network.

This approach to redefining the routing variables assumes that if a link fails the remaining links' traffic will be scaled up by the fraction of their operational time.

From the previous chapter on time delays we found a recursive equation for the average message delay from node i to node k as:

$$T_{ik} = \sum_{(i,j) \in E} \Phi_{ij}(k) [T_{jk} + \tau_{ij}] \quad (2.9)$$

Therefore if we can determine the routing variables for an unreliable network, then we can also determine the average message delay for an unreliable network.

$$\Theta'_{ij} = \theta_i \sum_k \Phi'_{ij}(k) \quad (2.10)$$

$$\Theta'_{ij} = \frac{\Theta_{ij}}{\pi_{ij}} \quad (2.11)$$

Since Θ'_{ij} corresponds to routing variables $\Phi'_{ij}(k)$, it must be scaled up by a factor of π_{ij} because the link is only operational a fraction π_{ij} of the time.

2.5. Time Delay

Let C_{ij} be the capacity of link (i,j) , then the average message delay on link (i,j) is denoted by τ_{ij} and is given by:

$$\tau_{ij} = \frac{1}{\mu C_{ij} - \Theta_{ij}} \quad (2.12)$$

The average delay, T for messages in the system is:

$$T = \frac{1}{P} \sum_{(i,j) \in E} \Theta_{ij} \tau_{ij} \quad (2.13)$$

Where E is the set of all links in the network, and P is the total message generation rate given by:

$$P = \sum_{i=1}^N \delta_i \quad (2.14)$$

The average message delay going from node i to node k is denoted by T_{ik} , and is given by:

$$T_{ik} = \begin{cases} \tau_{ij} + T_{jk} & \text{probability } \Phi_{ij}(k) \\ \tau_{ik} & \text{probability } \Phi_{ik}(k) \end{cases} \quad (2.15)$$

In other words we can state a recursive equation for the average message delay from node i to node k as:

$$T_{ik} = \sum_{(i,j) \in E} \Phi_{ij}(k) [T_{jk} + \tau_{ij}] \quad (2.16)$$

With the last equation we can determine the average message delay given the routing variables $\Phi_{ij}(k)$.

2.5.1. Calculation of Time Delay without Priorities

Let $\Xi_{ij}(k)$ be the total delay of a message sent to node k from node i routed over link (i,j) . The total delay is due to four parts.

1. T_{ik} = the time delay from node i to node k routed over some set of paths
2. ω_{ik} = the time delay due to queueing or the waiting time at node k before being processed

3. ϵ_{ik} = the time delay due to the service time

4. T_{ki} = the time delay from node k to node i routed over some set of paths for confirmation.

Then the following equation holds for total message delay:

$$\Xi_{ij}(k) = T_{ik} + \omega_{ik} + \epsilon_{ik} + T_{ki} \quad (2.17)$$

where we have the following equalities:

$$1) \quad \omega_{ik} = \begin{cases} \frac{\psi_i}{\mu^2} & \text{for queries} \\ \frac{\psi_i}{1 - \frac{\psi_i}{\mu}} & \text{for updates} \end{cases} \quad (2.18)$$

$$2) \quad \epsilon_{ik} = \begin{cases} 0 & \text{for queries} \\ 1/\mu & \text{for updates} \end{cases} \quad (2.19)$$

3) The values for T_{ij} and T_{ji} are set in the recursive equation

The time delay has been calculated as the time delay for an M/G/1 queue with variable service rates. In actuality this can be modeled as an M/G/1 queue with priority queueing and variable service rates. Using a good queueing discipline variable service rates should cause priority queueing. In priority queueing the messages that take less time to service should not have to wait in the queue for messages that take more time to service unless they are already in the process of being serviced.

2.5.2. Calculation of Time Delay with Priorities

In this section we consider the realistic case in which commanders have different rank and desire different files. With out loss of generality let us consider the one file case. In our application we have L commanders and S sensors. The commanders have different rank and different commands rate the importance of particular

files differently. Therefore the importance of commanders and importance of files to commanders can be combined into one general level of importance. Combining this with the importance of files to sensors given a list of $2L + S$ priorities. The reason that there are $2L + S$ priorities is because commanders can query and update files while sensors only update. commanders can request data on files that sensors update or command files. Commanders can also update command files. Sensors only update data stored on files. there are three service rates, one for commanders query $1/\eta$, one for commander update $1/\mu$, and one for sensor updates $1/\zeta$. Since the service times for queries are assumed to be negligible the queries always have the highest priorities. the service time for commanders updates is much less than that for sensor updates since changing a command takes less time than changing data. Therefore the commander updates have higher priorities than sensor updates. The arrival rates of commander queries to node i is λ_i , the arrival rates of commander updates to node i is ψ_i , and the arrival rates of sensor updates is α_i . given the $2L + S$ priority classes we would like to find the average waiting times for the different classes. This result is found in Kleinrock's [1.2] book "Queueing Systems II" and is calculated below. The waiting times for any particular class is given by Ω_p :

$$\Omega_p = \frac{\sum_{i=1}^{2L+S} \frac{\lambda_i \bar{X}_i^2}{2}}{(1 - \sum_{i=p}^{2L+S} \lambda_i \bar{X}_i)(1 - \sum_{i=p+1}^{2L+S} \lambda_i \bar{X}_i)} \quad (2.20)$$

Using the above equation we can find the average waiting times in the queue for our $2L + S$ priority classes. Where λ_i is the arrival rate of the i th priority class; \bar{X}_i is the average service time of the i th priority class; and \bar{X}_i^2 is the second moment for the service time of the i th priority class. The waiting times for our $2L + S$ priority classes can be broken down into six different cases. The six cases are as follows: $0 \leq p \leq L - 1$; $p = L$; $L + 1 \leq p \leq 2L - 1$; $p = 2L$; $2L \leq p \leq 2L + S - 1$; and $p = 2L + S$.

2.5.2.1. Case 1

This is the case of $0 \leq p \leq L - 1$:

$$\Omega_p = \frac{\sum_{i=1}^L \frac{\lambda_i \bar{X}_i^2}{2} + \sum_{i=L+1}^{2L} \frac{\psi_i \bar{X}_i^2}{2} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i \bar{X}_i^2}{2}}{(1 - \sum_{i=p}^L \lambda_i \bar{X}_i - \sum_{i=L+1}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)(1 - \sum_{i=p+1}^L \lambda_i \bar{X}_i - \sum_{i=L+1}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)} \quad (2.21)$$

The service rates for queries is zero so \bar{X}_i^2 and \bar{X}_i are zero for $1 \leq i \leq L$:

$$= \frac{\sum_{i=L+1}^{2L} \frac{\psi_i \bar{X}_i^2}{2} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i \bar{X}_i^2}{2}}{(1 - \sum_{i=L+1}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)(1 - \sum_{i=L+1}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)} \quad (2.22)$$

$$= \frac{\sum_{i=L+1}^{2L} \frac{\psi_i}{\mu^2} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta^2}}{(1 - \sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} - \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta})^2} \quad (2.23)$$

2.5.2.2. Case 2

This is the case of $p = L$:

$$\Omega_p = \frac{\sum_{i=1}^L \frac{\lambda_i \bar{X}_i^2}{2} + \sum_{i=L+1}^{2L} \frac{\psi_i \bar{X}_i^2}{2} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i \bar{X}_i^2}{2}}{(1 - \sum_{i=L+1}^L \lambda_i \bar{X}_i - \sum_{i=L+1}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)(1 - \sum_{i=L+1}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)} \quad (2.24)$$

Once again we use the fact that the service rates for queries is zero:

$$= \frac{\sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta}}{(1 - \sum_{i=L+1}^{2L} \psi_i \frac{1}{\mu} - \sum_{i=2L+1}^{2L+S} \alpha_i \frac{1}{\zeta})(1 - \sum_{i=L+1}^{2L} \psi_i \frac{1}{\mu} - \sum_{i=2L+1}^{2L+S} \alpha_i \frac{1}{\zeta})} \quad (2.25)$$

$$= \frac{\sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta}}{(1 - \sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} - \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta})^2} \quad (2.26)$$

2.5.2.3. Case 3

This is the case where $L+1 \leq p \leq 2L-1$:

$$\Omega_p = \frac{\sum_{i=1}^L \frac{\lambda_i \bar{X}_i^2}{2} + \sum_{i=L+1}^{2L} \frac{\psi_i \bar{X}_i^2}{2} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i \bar{X}_i^2}{2}}{(1 - \sum_{i=p}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)(1 - \sum_{i=p+1}^{2L} \psi_i \bar{X}_i - \sum_{i=2L+1}^{2L+S} \alpha_i \bar{X}_i)} \quad (2.27)$$

Once again we use the fact that the service rate for queries is zero:

$$= \frac{\sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta}}{(1 - \sum_{i=p}^{2L} \frac{\psi_i}{\mu} - \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta})(1 - \sum_{i=p+1}^{2L} \frac{\psi_i}{\mu} - \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta})} \quad (2.28)$$

2.5.2.4. Case 4

This is the case where $p=2L$, and we use the fact that the service rate for queries is zero:

$$= \frac{\sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta}}{(1 - \frac{\psi_{2L}}{\mu} - \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta})(1 - \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta})} \quad (2.29)$$

2.5.2.5. Case 5

This is the case where $2L \leq p \leq 2L+S-1$, and we use the fact that the service rate for queries is zero:

$$\Omega_p = \frac{\sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta}}{(1 - \sum_{i=p}^{2L+S} \frac{\alpha_i}{\zeta})(1 - \sum_{i=p+1}^{2L+S} \frac{\alpha_i}{\zeta})} \quad (2.30)$$

2.5.2.6. Case 6

This is the case where $p = 2L+S$, and we use the fact that the service rate for queries is zero:

$$\Omega_p = \frac{\sum_{i=L+1}^{2L} \frac{\psi_i}{\mu} + \sum_{i=2L+1}^{2L+S} \frac{\alpha_i}{\zeta}}{(1 - \frac{\alpha_{2L+S}}{\zeta})} \quad (2.31)$$

2.6. Capacity

What is desired is that flow of communication through the network is restricted in the sense that not all communication goes through one or two nodes. In other words we restrict the amount of flow over a link during a period of time. This flow can not be increased over capacity. The throughput is therefore limited over links so that more links are used.

3. ACCESSIBILITY

This section deals with the problem of finding the probability of accessibility of a distributed data base system in a failing computer network. Given the locations of the files stored in the system along with the probability of link and node failures, one can determine the accessibility of a user to any particular file in the system. A user is stationed at one of the nodes of the network, while any particular file may have multiple copies stored throughout the network. Starting with the link and node failure probabilities one can find the probabilities of inter-node accessibilities, p_{ij} 's. An inter-node accessibility, p_{ij} is simply the probability that node i can access whatever is at node j . Once the inter-node accessibilities are found, one can write a simple form for the accessibility of a user i to a file l in the system, $P_{il}(I_l)$, in terms of inter-node accessibilities and file allocation variables, x_{il} . $P_{il}(I_l)$ is the probability of accessibility of user i to file l , given the file allocation for file l is I_l . File allocation variables, x_{il} , are zero-one stating whether or not file l is stored at node i . If file l is stored at node i , then x_{il} is one; on the other hand, if file l is not stored at node i , then x_{il} is zero. The purpose of finding a simple form for $P_{il}(I_l)$ is to simplify the task optimal file allocation optimization.

3.1. Introduction

In the past researchers [3.3] have only dealt with the inter-node accessibilities. This is because one simply wanted to find the accessibility between a user and a service node; however, in the context of optimal file allocation in a distributed data base there may be more than one service node. In optimal file allocation in a distributed data base, it is necessary to determine accessibility between users and their desired files where there may be multiple copies of the desired files. Mahmoud and Riordon [3.1], and Laning and Leonard [3.2] have only dealt with approximations of accessibility in their optimal file allocation optimizations. They only have accessibility issue in the constraints of their optimal file allocation optimizations. In this way, their formulation only rejects certain file allocations because of the low level of accessibility; however, they do not try to maximize the accessibility between the users and the files. We propose to put the accessibility issue in the objective function of our optimal file allocation optimization; therefore, we will not only minimize the cost of a particular file allocation, but also maximize the probability of accessibility of user to files. In our formulation we will use the exact probability of accessibility. Once the probability of accessibility is known in terms of the allocation variables, x_{il} , and the inter-node accessibilities, p_{ij} , it may be put in the objective function. Since this term will be non-linear in the allocation variables, it is desired to find a simple form of the probability of accessibility so that one may easily linearize the objective function. First one must find the inter-node accessibilities.

3.2. Inter-Node Accessibilities

The term p_{ij} represents the probability of accessing node j from node i and is called the probability of inter-node accessibility. Many other researchers have dealt with the problem of inter-node accessibilities. Hansler, McAuliffe,

and Wilkov [3.3] have found an exact method of finding there inter-node accessibility. Their algorithm first finds a collectively exhaustive list of primal cutsets of the network. A primal cutset are a set of links when removed denies accessibility between the desired nodes. The next step is to transform the list to a mutually exclusive and collectively exhaustive list of primal cutsets. Then the sum of the probabilities of failure of each primal cutset is the probability of non-accessibility between the nodes. This is the complement of the inter-node accessibility for any two nodes. This may be extended to more than two nodes, since the probability of accessibility between three nodes can not be easily derived from the probability of accessibility between two nodes due to the possibility of dependence of link failures. Hansler, McAuliffe, and Wilkov [3.3] state that although the algorithm could take exponential time to evaluate, he states that the algorithm can handle networks as large as the Arpanet.

3.3. Accessibility

This section deals with the problem of finding the probability of accessibility of a distributed data base systems in a failing network. Given the location of the files stored in the system, we can determine the accessibility of a user from any particular node to a particular file in the system. Each file may be stored a multiple number of times. We present a simple form for the accessibility of a user to a file in the system.

Let us consider a distributed data base with only one file stored a multiple number of times. For our problem we can consider the one file system without loss of generality since we are only concerned with the accessibility of a file to a user. We wish to find $P_i(I)$ the probability of user i accessing the file with a particular file allocation I . This probability may be written as follows:

$$P_i(I) = \sum_{j=1}^N \prod_{\substack{k=1 \\ k \neq j}}^N (1 - x_k) x_j p_{ij} \\ + \sum_{h=1}^N \sum_{\substack{j=1 \\ j \neq h}}^N \prod_{\substack{k=1 \\ k \neq j, h}}^N (1 - x_k) x_j x_h [p_{ij} + (1 - p_{ij})p_{ih}] \dots \\ + \sum_{a=1}^N \sum_{\substack{j=1 \\ j \neq a}}^N \dots \sum_{\substack{j=1 \\ j \neq a, \dots}}^N \prod_{\substack{k=1 \\ k \neq j, a, \dots}}^N (1 - x_k) x_j x_a \dots x_{a_{n-1}} \\ \times [p_{ij} + (1 - p_{ij})p_{ih} + (1 - p_{ij})(1 - p_{ih})p_{ig} \dots] \quad (3.1)$$

Where x_i is a zero-one variable indicating whether the file is stored at node i or not. If the value of x_i is one the file is stored at node i ; on the other hand, if the value is zero the file is not stored at node i . p_{ij} is the probability of accessibility between nodes i and j .

The first term is the probability of accessibility of user i to the file stored at node j given the file is only stored at node j . The second term is the probability of accessibility of user i to the file given the file is stored only at nodes j and h . This probability is simply the probability that user i can access the file from node j or the probability that user i can not access the file at node j but can access the file at node h . Continuing in this fashion we arrive at the last term which states that the file is stored at all the nodes, so the probability of accessibility of user i to the file then becomes the probability that user i can access the file from node j ; or the probability that user i can access the file from node h but not from node j ; or the probability that user i can access the file from node g but not from node j and node h ; or etc..

We wish to rewrite this probability in a more simplified form. A more simplified form is one that requires fewer multiplications and the first, second, third, etc., and higher order terms are separated into different terms. One would like the first, second, third, etc., and higher order terms to be separated into different terms because then it is simple to linearize the zero-one integer program that is required for optimal file allocation, no matter if $P_i(I)$ is in the constraint or in the objective function. The simplified form will enable one to see any special structure in the equation. The form that we desire is recursive in nature, in the sense that the probability of accessibility for N nodes is just a simple relation from the probability of accessibility for $N-1$ nodes. The form we desire is the following:

$$P_i(I) = \sum_{j=1}^N x_j p_{ij} - \dots + (-1)^{N+1} \sum_{j=1}^N \sum_{\substack{h=1 \\ h \neq j}}^N \prod_{\substack{m=1 \\ m \neq j, h}}^N x_m p_{im} + (-1)^N \sum_{j=1}^N \prod_{\substack{m=1 \\ m \neq j}}^N x_m p_{im} + (-1)^{N+1} \prod_{m=1}^N x_m p_{im} \quad (3.2)$$

where p_{ij} is the probability of accessibility between nodes i and j . Let the probability of accessibility between user i and the file for an N node network be $P(N)$. Then using the above formula we may rewrite $P(N)$ as:

$$P(N) = x_N p_{iN} [1 - P(N-1)] \quad (3.3)$$

The proof of the equivalence is done in the proof of the theorem in the appendix. The theorem is simply the equivalence of the two forms of the probability of accessibility.

With this form of the probability of accessibility it is easy to see the impact of adding another node to the system.

3.4. Significance Of Results

We have been able to find a simple and exact form for the probability of accessibility between a user and a file that we are trying to allocate in our optimal file allocation optimization. A proof of the equivalence of the simple form and the definition of the probability of accessibility

is given. The probability of accessibility can be found in terms of the probability of inter-node accessibilities and the allocation variables. In turn the probability of inter-node accessibilities can be found in terms of the node and link failure probabilities. Therefore we can find an expression for the probability of accessibility in terms of the node and link failure probabilities and the allocation variables. This probability of accessibility can be placed in our objective function of our optimal file allocation optimization to not only minimize costs but also to maximize the probabilities of accessibility. Once this is done the non-linear integer program to solve our optimal file allocation can be linearized to a linear zero-one program. The problem now remains to develop a fast algorithm to solve our linear zero-one program for optimal file allocation of a distributed data base in an unreliable network using exact probabilities of accessibility.

3.5. Non-Accessibility Costs

There are two types of non-accessibility costs. There is the commander's non-accessibility costs and the sensor's non-accessibility costs. The commander's non-accessibility cost is due to the commander's desire for information stored in the Distributed Data Base, there is a cost associated with the commanders not being able to access the information due to link failures. This models the commander's accessibility costs. In the case of the sensors, multiple sensors may be feeding data into a particular file (FUSION). This data may be unusable if it is separate. Therefore, the data must be fused together. This may not be possible due to link failures. There is a cost associated with a sensor not being able to update the needed file. This models the sensor's files accessibility costs. For example, if one sensor is tracking an object and another sensor is trying to identify the object, then it is clear the sensor's data individually may mean nothing, but combined they can be very important.

Let us define the non-accessibility costs that we are trying to minimize in our formulation. Let us define the non-accessibility cost as the negative value of the probability of accessibility weighted by the importance of the files, commanders and (or) sensors. This has an averaging effect. In other words the formulation will try to allocate the files such that the more important the file the higher the probability of accessibility. Let us therefore try to write down some equations relating these facts to the non-accessibility of commanders and sensors.

3.5.1. Commanders

Let us define what are our unaccessability costs for commanders.

We define:

$$I_i(R(N_i)) = \beta_{li} A_i(I) \quad (3.4)$$

which denotes the accessibility of file l to the commander at node i weighted by the importance of file l to the commander at node i , β_{li} . $R(N_i)$ is the set of nodes accessible to node i . β_{li} is the importance of file l to commander i . $A_i(I)$ is a zero-one variable denoting whether the file l is accessible to node i ; therefore, $A_i(I)$ is one if the file is accessible and zero if the file is inaccessible. A file

is inaccessible either if there is not a path from node i to a node which file l is stored or if all the nodes containing file l are destroyed. Let us define the unaccessibility costs for commanders as:

$$-E \left[\sum_{i=1}^L \alpha_i I_l(R(N_i)) \right] \quad (3.5)$$

In other words the unaccessibility costs is the negative expected value of accessibility weighted by the importance of the commander α_i and the importance of the file to the commander.

If we now examine the last term in the minimization, we can simplify the expression. The expected value may be brought inside the summation. Since the importance of the commander i and the importance of file l to commander i are not probabilistic, we can simply take the expected value of the accessibility. However the expected value of the accessibility is simply the probability that commander i can access file l given the allocation of redundant copies of file l in the network. We have:

$$\begin{aligned} E \left[\sum_{i=1}^L \alpha_i I_l(R(N_i)) \right] &= E \left[\sum_{i=1}^L \alpha_i \beta_{li} A_i(l) \right]; \\ &= \sum_{i=1}^L \alpha_i \beta_{li} E[A_i(l)]; \\ &= \sum_{i=1}^L \alpha_i \beta_{li} P_{li}(l). \end{aligned} \quad (3.6)$$

where

$$P_{li}(l) = \begin{cases} 1 & \text{if } \Pr(\exists k \text{ s.t. } l \text{ at } N_k \in R(N_i)) \\ 0 & \text{if } \Pr(\nexists k \text{ s.t. } l \text{ at } N_k \in R(N_i)) \\ 0 & \text{if } \Pr(\forall k \text{ s.t. } l \text{ at } N_k, N_k \text{ destroyed}) \end{cases} \quad (3.7)$$

Where $P_{li}(l)$ for one file l , was found in the previous section to be:

$$\begin{aligned} P_i(l) &= \sum_{j=1}^N \prod_{k=1}^N (1 - x_k) x_j p_{ij} \\ &+ \sum_{h=1}^N \sum_{j=1}^N \prod_{k=1}^N (1 - x_k) x_j x_h [p_{ij} + (1 - p_{ij}) p_{ih}] \dots \\ &+ \sum_{a=1}^N \sum_{j=1}^N \dots \sum_{h=1}^N \prod_{k=1}^N (1 - x_k) x_j x_h \dots x_h [p_{ij} + (1 - p_{ij}) p_{ih} + (1 - p_{ij})(1 - p_{ih}) p_{ih} \dots] \end{aligned} \quad (3.8)$$

which simplifies to the following:

$$\begin{aligned} P_i(l) &= \sum_{j=1}^N x_j p_{ij} - \dots \\ &+ (-1)^{N+1} \sum_{j=1}^N \sum_{k=1}^N \prod_{m=1}^N x_m p_{im} \\ &+ (-1)^N \sum_{j=1}^N \prod_{m=1}^N x_m p_{im} + (-1)^{N+1} \prod_{m=1}^N x_m p_{im} \end{aligned} \quad (3.9)$$

where p_{ij} is the probability of accessibility between nodes i and j . Substituting the above equation into the non-accessibility cost of the commander gives us the final expression for the non-accessibility costs for commanders.

3.5.2. Sensors

The question is how do the sensors effect our problem of optimal file allocation in an unreliable network. The problem is that different sensors must fuse their data in data files so that their information becomes useful; whereas the data from the sensors individually is not as useful. There turns out to be a large difference if the data in the different sensors are fused or not fused. This is due to military applications where one sensor may be tracking position while the other sensor may be determining identity. In this case it is very important to match this data together, since if one does not know the identity of the object that one is tracking he does not know what action to take.

Our problem is how to model this aspect of sensors. This can be modeled in a similar way as we modeled commanders. There will be an importance of a file to a sensor. Then to model the fact that the sensor data must be fused together we assign an accessibility variable, which is non-zero only if the file is accessible to both sensors. There are S sensors in the system then the term that is to be added to the formulation is of the following form:

$$\begin{aligned} E \left[\sum_{i=1}^S \sum_{j=1}^S \alpha_i I_l(R(N_i) \cap R(N_j)) \right] &= E \left[\sum_{i=1}^S \sum_{j=1}^S U_{lij} A_{ij}(l) \right]; \\ &= \sum_{i=1}^S \sum_{j=1}^S U_{lij} E[A_{ij}(l)]; \\ &= \sum_{i=1}^S \sum_{j=1}^S U_{lij} P_{ij}(l). \end{aligned} \quad (3.10)$$

where U_{lij} is the usefulness of file l when the data from sensors i and j are fused there; and $A_{ij}(l)$ is defined below:

$$A_{ij}(l) = \begin{cases} 1 & \text{if } \exists k \text{ s.t. } l \text{ at } N_k \in R(N_i) \cap R(N_j) \\ 0 & \text{if } \nexists k \text{ s.t. } l \text{ at } N_k \in R(N_i) \cap R(N_j) \\ 0 & \text{if } \forall k \text{ s.t. } l \text{ at } N_k, N_k \text{ destroyed} \end{cases} \quad (3.11)$$

Where $P_{ij}(l)$ for one file l is by definition:

$$P_{ijl}(I_l) =$$

$$\begin{aligned} & \sum_{k=1}^N x_k p_{ijk} - \dots \\ & + (-1)^{N+1} \sum_{k=1}^N \sum_{m=1}^N \prod_{\substack{n=1 \\ n \neq k \\ n \neq m}}^N x_n p_{ijn} \\ & + (-1)^N \sum_{k=1}^N \prod_{m=1, m \neq k}^N x_m p_{ijm} + (-1)^{N+1} \prod_{m=1}^N x_m p_{ijm} \end{aligned} \quad (3.12)$$

where p_{ijk} is the probability of accessibility between nodes i and k , and nodes j and k . The problem with this new formulation is that it is extremely difficult to find the p_{ijk} 's since p_{ijm} can not be found easily from p_{ik} and p_{jk} since these two probabilities are not independent. These two probabilities are not independent because the path taken from node i to node k may cross the path taken from node j to node k . The final reformulation is as follows:

$$\min_{I_l} \sum_{l=1}^M C(I_l) =$$

$$\min_{I_l} \sum_{l=1}^M \left[\sum_{j=1}^N \left[\sum_{k \in I_l} \psi_{jl} \Xi_{jk} + \min_{k \in I_l} \lambda_{jl} \zeta_{jk} \right] + \sum_{k \in I_l} \sigma_{kl} - \sum_{i=1}^L \alpha_i \beta_{il} P_{il}(I_l) - \sum_{i=1}^S \sum_{j=1}^S U_{ij} P_{ijl}(I_l) \right] \quad (3.13)$$

4. NEW THEOREMS

The new theorems presented here along with theorems III and IV of last year's paper [4.1] show that the objective function is convex in the number of redundant file copies. Using this result, we may find an efficient heuristic algorithm to solve the optimal file allocation problem by an efficient search for the optimal number of redundant file copies.

4.1. Theorem XIII

If

$$C(I \sim [k]) \geq C(I \sim [1, 2]) \quad \text{for } k = 1, 2. \quad (4.1)$$

then

$$C(I) \geq C(I \sim [k]) \quad \text{for } k = 1, 2. \quad (4.2)$$

4.2. Theorem XIII in words and its implications

Theorem XIII states that on our cost graph if two descendants of a given vertex have a cost greater than the cost of the vertex, then the cost of the descendants of the two vertices is greater than either of the two vertices.

4.3. Theorem IV

Given an index set $X \supseteq I$, containing r elements with the following property:

$$C(I) \geq C(I \sim [x]) \quad \text{for each } x \in X. \quad (4.3)$$

Then for every sequence $R^{(1)}, R^{(2)}, \dots, R^{(r)}$, which are subsets of X , such that $R^{(k)}$ has k elements and $R^{(k)} \supset R^{(k+1)}$, the following is true:

$$C(I) \geq C(I \sim R^{(1)}) \geq C(I \sim R^{(2)}) \geq C(I \sim R^{(3)}) \geq \dots \geq C(I \sim R^{(r)}) \quad (4.4)$$

4.4. Theorem IV in words and its implications

Theorem IV states that if a given vertex has a cost greater than the cost of any vertex along the path leading from it, then the sequence of costs encountered along any one of these paths increase monotonically. Thus in order to find the optimal allocation policy, it is sufficient to follow every path of the cost graph in the reverse direction until the cost decreases and no further. This will give an allocation of local optimum of which the global optimum is one of them.

This allows us to reduce the solution space of the integer program. Once we find a local optimum then we know that any more file allocation is not required so that the integer program will not have to search for solutions in that part of the solution space.

4.5. Theorem XV

As the number of redundant copies of a file increases the objective function is convex. In other words the objective function first decreases to a local minimum, then increases as the number of redundant file copies increase.

4.6. Theorem XV in words and its implications

Theorem XV simply states that the objective function is convex in the number of redundant file copies. This will enable us to develop algorithms that will solve our formulation. The worst case of our algorithm will be better than any previous algorithm.

5. LINEARIZATION

$$\sum_{i=1}^M \min_{I_i} \left[\sum_{j=1}^N \left[\sum_{k=1}^N \psi_{ji} d_{jk} x_{ik} + \min_{k \in I_i} \lambda_{ji} d_{jk} \right] + \sum_{k=1}^N \sigma_{ki} x_{ik} - \sum_{i=1}^L \alpha_i \beta_{ii} P_{ii}(I_i) \right] \quad (5.1)$$

The problem with the formulation stated above is that there are two non-linear terms in the formulation. These terms will be seen to be non-linear in the x 's which are zero-one integers. Therefore our formulation as it currently stands is a non-linear zero-one integer program. However as it will soon be seen that a non-linear zero-one integer program can be changed to a linear zero-one integer program with the addition of some simple linear constraints. The two terms that are non-linear are listed below:

$$1) \min_{k \in I_i} \lambda_{ji} d_{jk} \quad (5.2)$$

$$2) \sum_{i=1}^L \alpha_i \beta_{ii} P_{ii}(I_i) \quad (5.3)$$

The non-linear portion of the above non-linear terms is simply:

$$1) \min_{k \in I_i} \lambda_{ji} d_{jk} \quad (5.4)$$

$$2) P_{ii}(I_i) \quad (5.5)$$

Now without loss of generality we can assume the case of only one file, so let $I = I_i$. With this our two non-linear terms are as follows:

$$1) \min_{k \in I} \lambda_j d_{jk} \quad (5.6)$$

$$2) \sum_{i=1}^L \alpha_i \beta_i P_i(I) \quad (5.7)$$

In the second equation the non-linear term is $P_i(I)$, which has been proven in theorem XII to be:

$$\begin{aligned} P_i(I) = & \sum_{j=1}^N x_j p_{ij} - \dots \\ & + (-1)^{N+1} \sum_{j=1}^N \sum_{i=1}^N \prod_{\substack{m=1 \\ m \neq j}}^N x_m p_{im} \\ & + (-1)^N \sum_{j=1}^N \prod_{\substack{m=1 \\ m \neq j}}^N x_m p_{im} + (-1)^{N+1} \prod_{m=1}^N x_m p_{im} \end{aligned} \quad (5.8)$$

Here clearly there are non-linear terms in each sum except the first sum. The non-linear terms are in x . Now let's see if we can also find an equivalent representation for the first non-linear term.

The first non-linear term is the minimization over all I of $\lambda_j d_{jk}$. However the terms λ_j and d_{jk} are simply the query traffic emanating from node j and the communication cost from node j to the file at node k respectively. These terms are known apriori, therefore it should be possible to represent the minimization as a series of non-linear terms. This is indeed the case. Without loss of generality let us suppose that we can reorder the product of these terms from smallest to the largest, with the smallest having the smaller index k . Then let us define the following:

$$\mu_{jk} = \lambda_j d_{jk} \quad (5.9)$$

Now assuming that the μ_{jk} 's are ordered in increasing order we can represent the minimization as a sum of non-linear terms as follows:

$$\min_{k \in I} \lambda_j d_{jk} =$$

$$\mu_{j1} x_1 + \mu_{j2} (1 - x_1) x_2 + \dots + \mu_{jN} \prod_{m=1}^{N-1} (1 - x_m) x_N \quad (5.10)$$

Now once again we see that these terms are once again non-linear in terms of the zero-one x 's of of integer program.

To reduce the above non-linear zero-one problem to a linear zero-one problem, we first consider the objective function. Let us define:

$$X_{ij \dots uv} = X_i X_j \dots X_u X_v \quad q = 2, \dots, Q \quad (5.11)$$

which takes value zero or one, where Q is the highest degree of nonlinearity. We then represent each nonlinear term in the objective function by terms of the above form and then examine its coefficient. If the coefficient of the non-linear term is positive, we introduce the following constraint equation:

$$X_i + X_j + \dots + X_u + X_v - q + 1 \leq X_{ij \dots uv} \quad (5.12)$$

If the coefficient of the non-linear term is negative, we introduce the following constraint equation:

$$X_i + X_j + \dots + X_u + X_v \geq q X_{ij \dots uv} \quad (5.13)$$

5. ALGORITHMS

5.1. Algorithm I

Our objective function has been shown to be convex in the number of redundant file copies in theorem XV. Since our objective function is convex, one can do a search on the number of redundant file copies using the golden ratio to find the optimal solution, while using the linear programming integer solutions from the reliable network solution as a starting point. Let N be the number of nodes in the network.

The steps of the algorithm are as follows:

1. Set the maximum number of file copies to r , from theorem I and II. Find the cost for r optimally placed copies with cost I_R
2. Eliminate all Q nodes that cost too much
3. Automatically add M nodes that cost so little so as to always be profitable
4. Find the linear programming solution to the underlying reliable network. The number of integer file allocations to the linear program is P . Use these P file allocations as an initial file allocation. Find the cost of the optimal file allocation for P redundant copies. This cost is I_P
5. Find the optimal file allocations for

$$A_1 = P + \text{integer} \left\lfloor \frac{(r-P)(\sqrt{5}-1)}{2} \right\rfloor \quad (6.1)$$

redundant file allocations with cost I_A . Find the optimal file allocations for

$$A_2 = P + \text{integer} \left\lfloor \frac{(r-P)(3-\sqrt{5})}{2} \right\rfloor \quad (6.2)$$

redundant file allocations with cost I_B .

6. If $A_1 = A_2$, then $A_1 = A_2$ is the optimal number of redundant file copies. Simply use the optimal file allocations for $A_1 = A_2$ redundant file copies.

7.

7.1. If $I_P = \max\{I_P, I_A, I_B, I_R\}$ go to step 8

7.2. If $I_R = \max\{I_P, I_A, I_B, I_R\}$ go to step 9

8. Set $P = A_1$ and $I_P = I_A$. Find the optimal file allocations for

$$A_1 = \frac{(r-P)(\sqrt{5}-1)}{2} \quad (6.3)$$

redundant file allocations, with cost I_A . Go to step 6.

9. Set $r = A_2$ and $I_R = I_B$. Find the optimal file allocations for

$$A_2 = \frac{(r-P)(3-\sqrt{5})}{2} \quad (6.4)$$

redundant file allocations, with cost I_B . Go to step 6.

Since the objective function is convex, the above algorithm must converge to the solution.

5.2. Algorithm II

For the totally reliable network let us examine the non-linear term:

$$\min_{x \in I} \lambda_j d_{j1} = \mu_{j1} x_1 + \mu_{j2} (1-x_1) x_2 + \dots + \mu_{jN} \prod_{m=1}^{N-1} (1-x_m) x_N \quad (6.5)$$

Let us redefine the variables:

$$\begin{aligned} x'_1 &= (1-x_1) \\ x'_2 &= (1-x_2) \\ &\vdots \\ x'_N &= (1-x_N) \end{aligned} \quad (6.6)$$

then the non-linear term becomes:

$$\mu_{j1} (1-x'_1) + \mu_{j2} x'_1 (1-x'_2) + \dots + \mu_{jN} \prod_{m=1}^{N-1} x'_m (1-x'_N) \quad (6.7)$$

$$= \mu_{j1} - \mu_{j1} x'_1 + \mu_{j2} x'_1 - \mu_{j2} x'_1 x'_2 + \dots + \mu_{jN} x'_1 x'_2 \dots x'_{N-1} - \mu_{jN} x'_1 x'_2 \dots x'_N \quad (6.8)$$

$$= \mu_{j1} + (\mu_{j2} - \mu_{j1}) x'_1 + (\mu_{j3} - \mu_{j2}) x'_1 x'_2 + \dots + (\mu_{jN} - \mu_{j(N-1)}) x'_1 x'_2 \dots x'_{N-1} - \mu_{jN} x'_1 x'_2 \dots x'_N \quad (6.9)$$

Now let's linearize the equation except the last term adding the appropriate constraints.

$$\mu_{j1} + (\mu_{j2} - \mu_{j1})y'_{11} + (\mu_{j3} - \mu_{j2})y'_{12} \dots + (\mu_{jN} - \mu_{j(N-1)})y'_{123 \dots N-1} \quad (6.10)$$

$$\begin{aligned} x'_1 + 1 - 1 &\leq y'_{11} \\ x'_1 + x'_2 + 2 - 1 &\leq y'_{12} \\ x'_1 + x'_2 + x'_3 + 3 - 1 &\leq y'_{123} \end{aligned} \quad (6.11)$$

$$x'_1 + x'_2 + \dots x'_{N-1} \leq y'_{123 \dots N-1}$$

The reason we can expand this in this manner is that $(\mu_{j(i+1)} - \mu_{ji})$ is always positive since they are ordered in ascending order. Therefore the coefficients of the non-linear terms are always positive and the non-linear terms may be linearized using the above constraints. The constraint matrix is.

$$\begin{pmatrix} 1 & 0 & 0 & \dots & 0 & -1 & 0 & 0 & \dots & 0 \\ 1 & 1 & 0 & \dots & 0 & 0 & -1 & 0 & \dots & 0 \\ 1 & 1 & 1 & \dots & 0 & 0 & 0 & -1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 1 & 1 & \dots & 1 & 0 & 0 & 0 & \dots & -1 \end{pmatrix} \quad (6.12)$$

We want to show that this matrix is unimodular. Then simply solve the problem in the x' 's and take the inverse to find the x 's. This problem can be solved in polynomial time since the constraint matrix is unimodular. It can also be solved using the simplex method which will yield only integer solutions. The simplex method needs to be calculated twice since the highest-order non-linear term has a negative coefficient and cannot be expanded in this manner. Therefore once we form the constraint equations that linearize the formulation, we solve the problem with the simplex method once with the highest-order non-linear term taking the value zero and once with the value one.

6.2.1. Steps

These are the steps of the algorithm:

1. Eliminate all Q nodes that cost too much
2. Automatically add M nodes that cost so little so as to always be profitable
3. Solve the problem with the simplex method y twice, once for the highest non-linear term zero and once with the the highest non-linear term one. The minimum of these two results is the optimal solution.

7. CONCLUSIONS

We have formulated the file allocation problem in a C^3 context where vulnerability is an issue. We consider the problem of optimally locating multiple copies of files in an unreliable computer network with significant time delays. Our formulation considers:

1. The probability of commander accessing files;
2. The probability of multiple sensors accessing files;
3. The importance of commanders;
4. The importance of sensors;
5. The importance of particular files to particular commanders.
6. Variable routing;
7. Priority queueing;
8. Time delays.

The theorems have provided ways to cut down on the possible file allocations (solution space) in which the integer program has to search. Therefore, we reduce the amount of time required to solve for a solution using integer programming.

We have extended and proved fifteen theorems, all applicable to the new formulation. These theorems reduce the search space of the integer program and also prove that the objective function is convex in the number of redundant file copies.

We have shown how to linearize the formulation to produce a linear zero-one integer programming formulation.

We have developed two fast algorithms to solve the problem of optimal file allocation in the reliable and unreliable networks. For the reliable network problem our algorithm finds a solution in polynomial time. This is a great advance over previous algorithms which all required heuristic non-polynomial time algorithms. For the unreliable network our heuristic algorithm takes advantage of the structure of the objective function of our problem. In the theorems presented, we have shown that the objective function is convex in the number of redundant file copies. Using this result we can simply use a search based on the golden ratio which will find the optimal number of redundant file copies. This algorithm will never need an exhaustive search, so it is superior to all integer programming solutions. This algorithm will also find an exact solution which is superior to all other heuristic algorithms that just find approximations to the solutions.

Further work will be concentrated on simulating these algorithms on a computer. We will also conduct sensitivity analysis on these results for networks up to fifteen nodes.

8. REFERENCES

- 1.1 S. Mahmoud, J.S. Riordon, "Optimal Allocation of Resources in Distributed Information Networks", ACM Transactions on Data Base Systems, Vol. 1, No. 1 (1976).
- 1.2 L. Kleinrock, "Stochastic Message Flow and Delay." Dover, N.Y., 1972.
- 1.3 L.J. Laning, M.S. Leonard, "File Allocation in a distributed Computer Communication Network", IEEE Trans. on Computers, vol. c-32, no. 3, March, 1983, pp. 232-244.

- 1 L. Kleinrock, "Queueing Systems II", Wiley, 1976.
- 2 V.O.K. Li, "Message Delay in a Communication Network with Unreliable Components", Proceedings of the 5th MIT(ONR) Workshop on C³ Systems, pp. 280-284.
- 1 S. Mahmoud, J.S. Riorden, "Optimal Allocation of Resources in Distributed Information Networks", ACM Transactions on Data Base Systems, Vol. 1, No. 1 (1976).
- 2 L.J. Laning, M.S. Leonard, "File Allocation in a distributed Computer Communication Network", IEEE Trans on Computers, vol. c-32, no. 3, March, 1983, pp. 232-244.
- 3.3 E. Hansler, G. K. Mcauliffe, and R. S. Wilkov, "Exact calculation of computer network reliability," in Proc. AFIPS 1972 Fall Nat. Comput. Conf., vol. 41, pp. 49-54.
- 4.1 M. H. Ma and M. Athans, "Optimal File Allocation Problems for Distributed Data Bases in Unreliable Computer Networks I", Proceedings of the 5th MIT(ONR) Workshop on C³ Systems, pp. 254-260.

A COLLISION RESOLUTION PROTOCOL WITH LIMITED CHANNEL SENSING - FINITELY MANY USERS

P. Papantoni-Kazakos, Glenn D. Marcus, and Michael Georgiopoulos

University of Connecticut
Storrs, Connecticut 06268

ABSTRACT

In this paper, we consider the random-access of a single slotted channel by a finite number of independent, data transmitting bursty users. We adopt the assumption that each user monitors the channel only while he is blocked. We also assume that the channel outcomes (visible to each user) are binary. That is, each channel slot is perceived as either a noncollision, or as a collision slot. We disregard propagation delays. For the above model, we propose and analyze a collision resolution protocol (BCRLS) with tree search characteristics. For identical users with binomial transmission processes, we find lower bounds on the BCRLS throughput, and we compute upper bounds on the induced delays in transmission. We compare our results with those induced by the dynamic tree protocol of Capetanakis; where the feedback sensing is continuous in the latter.

I. INTRODUCTION

We consider a number of data-transmitting bursty users who request access to a single network resource. We assume that the users do not communicate with each other directly, and that their data are formatted into packets of identical length. Such a user model arises, for example, when a number of computer terminals access a single host computer. Let us further assume that the network resource is a single transmission channel whose time is slotted. The length of each channel time slot is equal to one packet. Also, each user can attempt transmission of a packet, starting only at the beginning of a slot. Given the general model above, a variety of transmission protocols can be devised depending on the specific characteristics of the user and channel models. Such characteristics include finite number of well-identified users versus an asymptotically large number of ill-specified users, as well as various levels of feedback information provided to the users by the channel, [1-17].

In the present paper, we assume finitely many users, and we adopt similar user and channel models as in [4,6]. In contrast to those models as well as to the collision-free models, however, we assume that each user inspects the broadcast feedback only while he is blocked. By blocking we mean the existence of some unsuccessfully transmitted packet in the user's buffer. Our assumption is the same as in [14], and it eliminates the often undesirable requirement that all users monitor the channel constantly, even when empty. In contrast to [14], we assume that the feedback is binary, distinguishing between collision and noncollision slots. For the above model, we propose and analyze a collision resolution protocol. We name this protocol Binary Collision Resolution with Limited Sensing (BCRLS).

This work was supported jointly by the National Science Foundation under the Grant ECS-8119885, and the Air Force Office of Scientific Research under Grant AFOSR-78-3695.

II. THE BCRLS PROTOCOL

As in [6], we assume that 2^n identical, independent, and packet-transmitting users share a single slotted channel. If a single packet is transmitted within a slot, it is received correctly. If at least two packets are simultaneously transmitted, a collision occurs, all information contained in the collided packets is lost, and retransmission is necessary. The binary outcome of each slot (empty or busy with one packet, collision) is broadcasted to all users without propagation delays. Let the system start operating at time zero, and let time be measured in slot units. Initially, each user does not inspect the feedback, and he is free to transmit a packet in any slot. Let some user transmit his first packet at time t . He then inspects the feedback that corresponds to slot t . If he sees success, he stops inspecting the feedback until his next transmission. If, instead, he sees a collision, he initiates the BCRLS protocol for collision resolution, while inspecting the feedback continuously. The user perceives the collision as resolved, as soon as his collided packet is successfully transmitted. He then stops inspecting the feedback, until his next transmission. Transmission of new arrivals is not attempted by the user, until his own collision has been resolved. In this section, we will briefly describe the BCRLS protocol and we will analyze its operational characteristics. Since those aspects are independent of the packet-generating process per user, we will make no assumptions on the latter at this point. We will need such assumptions, only when we study the expected delays induced by the protocol, and its stability properties. More details on the BCRLS can be found in [18].

A. The BCRLS General Operation

Given 2^n users, consider the binary tree with 2^n leaves. The tree has $n+1$ levels of depth, numbered from 0 to n . Depths 0 and n correspond respectively to the root and the leaves of the tree. In general, there exist 2^i nodes at depth i . Each of these nodes is the root of a binary subtree with 2^{n-i} leaves. Each tree node beyond depth 0 is identified by a binary codeword, where the codeword of each node at depth $i: i \geq 1$ contains i bits. Consider the depths of the binary tree evolving sequentially from left to right (as in figure II.1), and let some node at depth $i: 1 \leq i \leq n-1$ be identified by the binary codeword $x_1 x_2 \dots x_i$. Then, the two nodes at depth $i+1$ that branch off node $x_1 x_2 \dots x_i$ are identified by the codewords $x_1 \dots x_i 0$ and $x_1 \dots x_i 1$; node $x_1 \dots x_i 1$ lies under node $x_1 \dots x_i 0$. The two nodes at depth 1 that branch off the tree root R are identified by the length one codewords 0 and 1, where node 1 lies under node 0. It is clear from the above that each binary codeword $x_1 \dots x_i: 1 \leq i \leq n$ identifies a single tree node at depth i ; thus, it also identifies the unique path that connects this node with the tree root. In particular, each one of the 2^n distinct binary codewords of length n identifies a single tree leaf. Consider now a one-to-one correspondence between the 2^n

users and the 2^n binary codewords of length n . Then, each encoded user is uniquely identified by a single binary codeword of length n . But, as we saw above, each such codeword also identifies a single leaf on the $(n+1)$ -depth binary tree. Thus, there exists a one-to-one correspondence between those tree leaves and the 2^n encoded users. Let the 2^n users be encoded by binary sequences of length n , and let each user have in its memory a reproduction of the $(n+1)$ -depth binary tree. Each user considers himself placed on the tree leaf whose codeword coincides with his own. If we number the 2^n users from 1 to 2^n , the codeword of the i th user is the length n binary representation of the number $i-1$. From now on we will identify some user either by his number or by his codeword. Let t_0 be some time instant such that all collisions involving packets from user i have been previously resolved. Let user i transmit again at time t , where $t \geq t_0$ and the user did not transmit in the time interval (t_0, t) . Let there be a collision in slot t . User i observes this collision and starts the motions for its resolution. In the process, he uses the $(n+1)$ -depth binary tree in his memory, inspecting continuously the feedback. He traces sequentially neighboring tree nodes, in an order dictated by the feedback. At the same time, he imagines himself placed at different tree nodes depending on the relationship between the current node in the trace and the user's codeword. The user transmits, when the node he has placed himself at coincides with the trace node. To explain the collision resolution process coherently, we need to discriminate between node tracing and user node self-placing. To do that, we first present the following definition.

Definition 1

At some point in time, user i is blocked if he is in the process of resolving a collision. The blocked user i is in resolution mode 2^k , $0 \leq k \leq n$ if he imagines himself placed at a node that lies at the tree depth $n-k$. If $k \leq n$, and if $x_1 x_2 \dots x_n$ is the user's codeword, the node's codeword is $x_1 x_2 \dots x_{n-k}$. The user is in resolution mode 2^k and active (2_a^k) if he is in resolution mode 2^k and transmits. He is in resolution mode 2^k and withholding (2_w^k) if he is in resolution mode 2^k and does not transmit. The blocked user i is in state $(2_a^k, y_1 \dots y_\ell)$ ($2_w^k, y_1 \dots y_\ell$), $\ell \leq n$ if he is in resolution mode 2^k , his node tracing has reached the tree node whose binary codeword is $y_1 \dots y_\ell$, and he is respectively active or withholding.

Let $x_1 \dots x_n$ be the codeword of user i . Let the user be unblocked at time $t-1$, and let him transmit and be blocked at time t . Then, the BCRLS protocol performed by the user is described by the following statements.

- At time t , the user imagines himself placed at the root of the tree and transmitting. Observing collision, he also starts his node-tracing at the tree root. Thus, at time t user i is in state $(2_a^n, R)$.
- Let at time t_1 user i be blocked and in resolution mode 2^k , $0 \leq k \leq n$. Then, by definition 1, the user has placed himself at node $x_1 \dots x_{n-k}$. If he is in resolution mode 2_a^k , he transmits in slot t_1 . If he is in resolution mode 2_w^k , he does not. He cannot be in resolution mode 2_a^k unless his node-tracing has reached node $x_1 \dots x_{n-k}$. Let the user be in resolution mode 2_a^k at time t_1 . Then,

- If he observes success in slot t_1 , he becomes unblocked.
- If he observes collision in slot t_1 , he moves at time t_1+1 to resolution mode:

$$2_a^{k-1}; \text{ if } x_{n-k+1} = 0, \text{ and } k \geq 1$$

$$2_w^{k-1}; \text{ if } x_{n-k+1} = 1, \text{ and } k \geq 1$$

$$2_a^{n-1}; \text{ if } x_1 = 0, \text{ and } k = 0$$

$$2_w^{n-1}; \text{ if } x_1 = 1, \text{ and } k = 0$$

Let the user be in resolution mode 2_w^k at time t_1 . Then,

- If he observes a collision in slot t_1 , he moves at time t_1+1 to resolution mode:

$$2_a^{n-1}; \text{ if } k = 0, \text{ and } x_1 = 0$$

$$2_w^{n-1}; \text{ if } k = 0, \text{ and } x_1 = 1$$

Otherwise, he remains in resolution mode 2_w^k .

- If he observes an either empty or a successful slot t_1 , he moves at time t_1+1 to resolution mode:

$$2_a^{k-1}; \text{ if his node-tracing reaches node } x_1 \dots x_{n-k+1} \text{ at time } t_1+1.$$

Otherwise, he remains in resolution mode 2_w^k .

- Let at time t_1 user i be blocked. Let $y_0 y_1 \dots y_\ell$; $0 \leq \ell \leq n$ be the node reached by the user's node-tracing at time t_1 ; where if $\ell=0$ the node is the tree root. Then,

- If the user observes a collision at t_1 , he moves at time t_1+1 his node-tracing to node:

$$y_1 \dots y_\ell 0; \text{ if } 1 \leq \ell \leq n-1$$

$$0; \text{ if either } \ell=0 \text{ or } \ell=n$$

- If the user is in withholding resolution mode at time t_1 , and observes either empty or successful slot t_1 , he moves at time t_1+1 his node-tracing to node:

$$y_1 \dots y_{m-1} 1; 1 \leq m \leq \ell$$

; where $m; y_j=1$; $m+1 \leq j \leq \ell$ and $y_m=0$, for $j \geq 1$ and $y_1 \dots y_\ell$ such that not all bits y_j are equal to one.

- If the user is in active resolution mode at time t_1 , and observes success, he becomes unblocked.

Statements 1 to 3 above, basically describe a tree search as in [6]. The difference here is that this search is not performed simultaneously by all users. Thus, it is possible that when some user's search reaches a leaf node, a collision occurs. Then, as statements 1 to 3 indicate, the user interprets this collision as a root collision, and he reinitiates his tree search.

B. Properties of the BCRLS

As we explained previously, the BCRLS tree search is not performed simultaneously by all the 2^n users. Indeed, each user initiates his own BCRLS when he becomes blocked. Therefore, at some point in time, there may be some blocked users at various levels of their tree search, and some unblocked users. The unblocked users do not inspect the feedback. Among the blocked users, there will be some in active resolution mode and some in withholding resolution mode. Any possible collision will be caused by the active users, and will be observed by all the blocked users. Among the users who are blocked and in withholding mode, we will single

out those who are either in state $(2^k, x_1 \dots x_{n-k-1} 0)$; $0 \leq k \leq n-1$, or in some state $(2^k, x_1 \dots x_{n-k-1} 01 \dots 1)$; $1 \leq k \leq n-1$, where $x_1 \dots x_{n-k-1}$ are the first $n-k$ bits of the user's codeword. To discriminate between those users and the remaining blocked and withholding users, we will denote the resolution mode of the first 2^k users. As shown in [18], if a user is in resolution mode 2^k and observes an either empty or a successful slot, he moves to resolution mode 2^k . Let us denote,

$(\{2^k, N_k\}; 0 \leq k \leq n, \{2^k, P_k\}; 0 \leq k \leq n-1, [M])$: The event that at some point in time there are M unblocked users, N_k ; $0 \leq k \leq n$ users in resolution mode 2^k , and P_k ; $0 \leq k \leq n-1$ users in resolution mode 2^k . (1)

We now present a proposition whose proof is in [18].

Proposition 1

Given 2^n users, consider the event $(\{2^k, N_k\}; 0 \leq k \leq n, \{2^k, P_k\}; 0 \leq k \leq n-1, [M])$, at some point in time. Then, for every $k \leq n-1$ such that $N_k > 0$, the first $n-k$ codeword bits of the N_k users are identical; thus the N_k users branch off the same tree node at depth $n-k$. The same is true for the P_k users, if $P_k > 0$. Therefore, $N_k \leq 2^k$, and $P_k \leq 2^k$.

Let us now suppose that at some time instant t , the event $(\{2^k, N_k\}; 0 \leq k \leq n, \{2^k, P_k\}; 0 \leq k \leq n-1, [M])$ occurs. For some k such that $0 \leq k \leq n$, let the codewords of the N_k users have the common prefix $x_1(k) \dots x_{n-k}(k)$; where for $k=n$ this prefix is the tree root. Among those users, let us have N_{k1} with common codeword prefix $x_1(k) \dots x_{n-k}(k)0$; let us have N_{k2} with common codeword prefix $x_1(k) \dots x_{n-k}(k)1$. For some k such that $0 \leq k \leq n-1$, let the codewords of the P_k users have the common prefix $y_1(k) \dots y_{n-k-1}(k)1$. Then, $y_1(k) \dots y_{n-k-1}(k)1$ cannot be identical to $x_1(k) \dots x_{n-k}(k)$. If $P_0 + N_0 = 2$, then $y_1(0) \dots y_{n-1}(0)$ and $x_1(0) \dots x_{n-1}(0)$ are identical. This last statement evolves from proposition 1, and it is proved

in [18]. If $\sum_{k=0}^n N_k = 0$, slot t is empty. If $\sum_{k=0}^n N_k = 1$,

slot t is a successful slot, and then the one active

user becomes unblocked. If $\sum_{k=0}^n N_k \geq 2$, slot t is a

collision slot. The above observations lead to some simplifications regarding the events in (1) and their transitions in time. To show that, let us first denote,

$(\{2^k, N_k\}; 1 \leq k \leq n, [M])$: The event such that at some point in time there are M unblocked users, N_k ; $1 \leq k \leq n$ users in resolution mode 2^k , and no users in resolution mode 2^0 . (2)

We now present a proposition whose proof is in [18].

Proposition 2

Let at some time t the collision event $(\{2^k, N_k\}; 1 \leq k \leq n, [M])$ occur. Let for some k ; $1 \leq k \leq n$ be N_{k1} users with codeword prefix $x_1(k) \dots x_{n-k}(k)0$, and N_{k2} users with codeword prefix $x_1(k) \dots x_{n-k}(k)1$. Then, none of the $\sum_{k=1}^n N_{k2}$ users and the users that are in withholding mode at time t become unblocked, before all the $\sum_{k=1}^n N_{k1}$ users do, and before all the users who become blocked in the mean time transmit successfully.

Let us now denote,

$L(\{2^k, N_k\}; 1 \leq k \leq n, [M])$: The expected number of slots needed for the resolution of the collision represented by the event $(\{2^k, N_k\}; 1 \leq k \leq n, [M])$ in (2), just after this collision has been observed. (3)

From the derivations in [18], we obtain,

$$L(\{2^k, N_k\}; 1 \leq k \leq n, [M]) = \sum_{k=1}^n N_k \geq 2$$

$$2 + L(\{2^k, N_k\}; 1 \leq k \leq n, [M]) = \sum_{k=1}^n N_k \geq 2$$

$$\text{if } m + \sum_{k=1}^n N_{k1} \leq 1$$

$$2 + L(\{2^k, N_k\}; 1 \leq k \leq n, [M]) = \sum_{k=1}^n N_k \geq 2$$

$$+ L(\{2^k, N_k\}; 1 \leq k \leq n, [M]) = \sum_{k=1}^n N_k \geq 2$$

$$\text{if } m + \sum_{k=1}^n N_{k1} \geq 2 \quad (4)$$

Let us now consider the N_k users who are in resolution mode 2^k . Those users have a common codeword prefix $x_1(k) \dots x_{n-k}(k)$, and they are at most 2^k . Let us define by $P(N_{k1}/N_k)$, the probability that given the number N_k , there are N_{k1} users whose codeword prefix is $x_1(k) \dots x_{n-k}(k)0$. Then, N_{k2} will be equal to $N_k - N_{k1}$, and clearly,

$$P(N_{k1}/N_k) = \frac{\binom{2^{k-1}}{N_{k1}} \binom{2^{k-1}}{N_k - N_{k1}}}{\binom{2^k}{N_k}}; \max(0, N_k - 2^{k-1}) \leq N_{k1} \leq \min(2^{k-1}, N_k) \quad (5)$$

III. THE BCRLS PERFORMANCE

As we saw in section II, if at some time instant t the event $(\{2^k, N_k\}; 1 \leq k \leq n, [M])$ occurs, then M denotes the number of users who are unblocked at t . To this point, we made no assumptions as to the transmission characteristics of those users. Here, we will assume, as in [6], that an unblocked user transmits with probability q per slot. Thus, if at some point in time the number of unblocked users is M , then the probability $Q(m, M)$ that m users will transmit is given by the following expression.

$$Q(m, M) = \binom{M}{m} q^m (1-q)^{M-m}; 0 \leq m \leq M \quad (6)$$

The implication behind the expression in (6) is that the unblocked users transmit independently. This is consistent with the independence assumption made at the beginning of this paper. Using expressions (4), (5), and (6), we can express an equation for the expected value $L(\{2^k, N_k\}; 1 \leq k \leq n, [M])$. This equation is given in [18], it relates the expected values of different events as in (2), and it determines a linear system of equations whose solution is the set of expected values as in (3).

Definition 2

Given 2^n users, the throughput $2^n q_n$ of the BCRLS protocol is such that every q less than q_n provides a bounded and nonnegative solution for the linear system in [18], and no q value larger than q_n does.

For given q value, the solution of the linear

system in [18] can be obtained numerically. The throughput $2^n q_n$ can be also found numerically through the trials of different q values. We will present such numerical results in section IV. Let us define,

$D(2^n, q)$: The expected number of slots for the resolution of some collision at $t+1$, given 2^n users, given that at t all users are unblocked, and given that each unblocked user transmits with probability q .

The quantity $D(2^n, q)$ above is parallel to the expected delay $E\{\text{delay}\}$ in [6], and it is clearly given by the following expression.

$$D(2^n, q) = 1 + \sum_{N=2}^{2^n} \frac{2^n}{N} q^N (1-q)^{2^n-N} L(\{2^n, N\}, [2^n-N]) \quad (7)$$

For 2^n users, the stability region of the BCRLS consists of those $2^n q$ values that provide a bounded and nonnegative solution for the linear system in [18]. For such $2^n q$ values, the lengths $L(\{2^n, N\}, [2^n-N])$; $2 \leq N \leq 2^n$ will be finite; thus collisions that are signified by events $(\{2^n, N\}, [2^n-N])$ will end in finite time. Furthermore, due to proposition 2, when such collisions end, all the users will be unblocked. In conclusions, the quantity $D(2^n, q)$ provides the expected length of independent episodes, and $D(2^n, q)-1$ bounds from above the expected per packet delays (measured in slots).

IV. NUMERICAL RESULTS AND CONCLUSIONS

Due to lack of space, we exhibit here numerical results, without discussion. In table IV.1, lower bounds on the BCRLS throughput are listed. In figure IV.1, upper bounds on transmission delays are included. In table IV.3, lower bounds $2^n q_{an}^*$ on the maintainable by the BCRLS arrival rates are exhibited, and compared with the corresponding such rates T_n and T_n^d in [6], for the nondynamic and the dynamic tree search respectively. Details on the above aspects can be found in [18].

# of Users	$2^n q_n^*$
$2^3; n=3$.47192
$2^4; n=4$.44336
$2^5; n=5$.43006
$2^6; n=6$.42368
$n \rightarrow \infty$.36

Table IV.1

Lower Bounds on the BCRLS Throughput

# of Users	$2^n q_{an}^*$	T_n	T_n^d
$2^3; n=3$.725
$2^4; n=4$.630
$2^5; n=5$.560
$2^6; n=6$.515	.420	.675
$n \rightarrow \infty$.36	.346	.429

Table IV.3

Lower Bounds on the Maintainable by the BCRLS Input Rates

REFERENCES

- [1] N. Abramson, "The ALOHA System-Another Alternative for Computer Communications," AFIPS Conf. Proc., Fall Joint Comp. Conf. Vol. 37, pp. 281-285, 1970.
- [2] L.G. Roberts, "ALOHA Packet System with and Without Slots and Capture," ASS. Note 8, ARPA Network Info. Ctr., Stanford Res. Inst., Stanford, CA., June 1972.
- [3] L. Kleinrock and F.A. Tobagi, "Packet Switching in Radio Channels: Part I - Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics," IEEE Trans. Comm., Vol. COM-23, pp. 1400-1416, Dec. 1975.
- [4] J.F. Hayes, "An Adaptive Technique for Local Distribution," IEEE Trans. Comm., Vol. COM-26, pp. 1178-1186, August 1978.
- [5] J.I. Capetanakis, "Tree Algorithms for Packet Broadcast Channels," IEEE Trans. Inf. Th., Vol. IT-25, pp. 505-515, Sept. 1979.
- [6] J.I. Capetanakis, "Generalized TDMA: The Multi-Accessing Tree Protocol," IEEE Trans. Comm., Vol. COM-27, pp. 1476-1484, Oct. 1979.
- [7] B.S. Tsybakov and V.A. Mikhailov, "Free Synchronous Access of Packets to a Broadcast Channel with Feedback," Probl. Peredachi Inf., Vol. 14, No. 4, pp. 32-59, 1978.
- [8] B.S. Tsybakov, "Resolution of a Conflict of Known Multiplicity," Probl. Peredachi Inf., Vol. 16, No. 2, pp. 69-82, 1980.
- [9] L. Georgiadis and P. Papantoni-Kazakos, "A Collision Resolution Protocol for Random Access Channels with Energy Detectors," IEEE Trans. Comm., Vol. COM-30, pp. 2413-2420, Nov. 1982.
- [10] L.W. Hansen and M. Schwartz, "An Assigned-Slot Listen-Before Transmission Protocol for A Multi-access Data Channel," IEEE Trans. Comm., Vol. COM-27, pp. 846-857, June 1979.
- [11] L. Kleinrock and M. Scholl, "Packet Switching in Radio Channels: New Conflict-Free Multiple Access Schemes for a Small Number of data Users," Proceedings ICC Conf., June 1977, pp. 22.1-105-22.1-111.
- [12] J.W. Mark, "Global Scheduling Approach to Conflict-Free Multiaccess Via a Data Bus," IEEE Trans. Comm., Vol. COM-26, pp. 1342-1352.
- [13] G.D. Marcus and P. Papantoni-Kazakos, "Dynamic Scheduling Protocols for a Multiple-Access Channel," IEEE Trans. Comm., submitted. Also, part of work in Proceedings, 1981 ICC, Denver, June.
- [14] B.S. Tsybakov and N.D. Vvedenskaya, "Random Multiple-Access Stack Algorithm," Probl. Peredachi Inf., Vol. 16, No. 3, pp. 80-94, 1980.
- [15] M.G. Hlnchyj and R. Gallager, "Multiaccess of a slotted channel by Finitely Many Users," Proceedings, NTC' 81, New Orleans, Nov., pp. D4.2.1-D4.2.7.
- [16] L. Merakos, "Limited Sensing Random Multiple Access Using Binary Feedback," 1983 Johns Hopkins Conference on Information Sciences and Systems.
- [17] P. Papantoni-Kazakos and G.D. Marcus, "Random Access with Limited Sensing," 21st IEEE Conference on Decision and Control, Dec. 1982.
- [18] P. Papantoni-Kazakos, G.D. Marcus and M. Georgiopoulos, "A Collision Resolution Protocol with Limited Channel Sensing-Finitely Many Users," Univ. of Conn., Tech. Rep. TR-83-2, EECs Dept., Feb. 1983.

List of Attendees of the 6th MIT/ONR Workshop on C³ Systems

Mark Alexandridis
Technical Staff Member
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Robert Anthony
Member of Technical Staff
The MITRE Corporation
Mail Stop W961
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6817

Michael Athans
Professor EECS
Massachusetts Institute of Technology
Room 35-406/LIDS
77 Massachusetts Avenue
Cambridge, MA 02139
Tel: (617) 253-6173

Alan Bamford
SCICON Ltd.
49 Berners Street
London W1P 4AQ
ENGLAND
Tel: 01-580-5599

Vernon P. Bailey
Electronics Engineer
Naval Underwater Systems Center
Code 3513 - Bldg. 1171
Newport, RI. 02841
Tel: (401) 841-2631

John S. Baras
Professor EE
University of Maryland
College Park, MD 20742
Tel: (301) 454-6876

Thomas A. Bartee
Technical Aid
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22107
Tel: (703) 448-7849

Dimitri P. Bertsekas
Professor EECS
Massachusetts Institute of Technology
Room 35-210/LIDS
77 Massachusetts Avenue
Cambridge, MA 02139
Tel: (617) 253-7267

Gilmer Blankenship
Professor EE
University of Maryland
College Park, MD 20742
Tel: (703) 941-6206

Alan L. Blitz
Technical Staff Member
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Paul L. Bongiovanni
Electronics Engineer
Naval Underwater Systems Center
Bldg. 1171-3
Newport, RI 02841
Tel: (401) 841-2613

Christopher Bowman
VERAC, Inc.
10975 Torreyana
San Diego, CA 92171
Tel: (619) 459-5550

Edward C. Brady
Technical Director
The MITRE Corporation
MS-W670
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6201

James V. Bronson
Colonel USMC
Marine Corps Tactical Systems
Support Activity
Marine Corps Base
Camp Pendleton, CA 92055
Tel: (619) 725-2618

Stuart L. Brodsky
11042 Seven Hill Lane
Potomac, MD 20854

Rex V. Brown
Decision Science Consortium, Inc.
7700 Leesburg Pike, Suite 421
Falls Church, VA 22043
Tel: (703) 790-0510

Nicholas J. Carriero
Research Team Leader
U. S. Army Human Engineering Lab.
Aberdeen Proving Ground, MD 21005
Tel: (301) 278-5879

Preceding page blank

David A. Castanon
Research Scientist
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Julian L. Center, Jr.
JAYCOR
300 Unicorn Park Drive
Woburn, MA 01801
Tel: (617) 933-5805

Tsu-Shuan Chang
Assistant Professor EE
SUNY at Stony Brook
Stony Brook, NY 11794
Tel: (516) 246-8612

C. Y. Chong
Advanced Information and Decision Systems
201 San Antonio Circle, Suite 286
Mountain View, CA 94022
Tel: (415) 941-3912

W. Peter Cherry
Vice President
Vector Research, Inc.
P. O. Box 1506
Ann Arbor, MI 48106
Tel: (313) 973-9210

Gerald A. Clapp
Program Manager
Naval Ocean Systems Center
Code 8105
San Diego, CA 92152
Tel: (619) 225-2044

J. Dallas Clarke, IV
Program Manager
Defense Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, VA 22209
Tel: (202) 694-1821

Marvin S. Cohen
Program Manager
The MAXIMA Corporation
7315 Wisconsin Avenue
Suite 900N
Bethesda, MD 20814
Tel: (301) 652-6366

Bruce P. Cooper
Asst. Manager
Hughes Aircraft Company
P. O. Box 3310-Bldg. 617/C210
Fullerton, CA 92634
Tel: (714) 732-5211

G. C. Corynen
Lawrence Livermore National Laboratory
P. O. Box 5504, L-156
Livermore, CA 94550
Tel: (415) 422-8537

Ronald J. Coulter, Lt. Col.
Branch Head, C³ Division (D103F)
Development Center, MCDEC
Quantico, VA 22134
Tel: (703) 640-3161

Harry E. Crisp
Head, Weapons Control Development Branch
Code N51
Naval Surface Weapons Center
Dahlgren, VA 22448
Tel: (703) 663-7431

J. B. Cruz, Jr.
Professor EE
Coordinated Science Laboratory
University of Illinois
1101 W. Springfield Avenue
Urbana, ILL. 61801
Tel: (217) 333-0230

William E. Daley
Member Technical Staff
Bell Telephone Laboratories
Room 4C-246A
Whippany Road
Whippany, NJ 07981

Paul K. Davis
Director, Strategy Assessment Center
The Rand Corporation
1700 Main Street
Santa Monica, CA 90406
Tel: (213) 393-0411

Francis C. Deckelman
Program Manager
Naval Electronic Systems Command
Code 6131
2511 Jefferson Davis Hwy.
Arlington, VA 22217
Tel: (202) 692-6090

James Deckert
Manager Decision Systems Section
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

J. R. Delaney
Member Technical Staff
MIT Lincoln Laboratories
P. O. Box 73
Lexington, MA 02173
Tel: (617) 863-5500 Ext. 2541

Hugh A. L. Dempsey
Team Leader
USA Human Engineering Laboratory
ATTN: DRXHE-FT
Aberdeen Proving Ground, MD 21005
Tel: (301) 278-5882

Lawrence J. Dooks
Member Technical Staff
Bell Telephone Laboratories
Dep't. 45311
Whippany Road
Whippany, NJ 07981
Tel: (201) 386-4832

Elizabeth Ducot
Technical Staff
Massachusetts Institute of Technology
Room 35-410/LIDS
77 Massachusetts Avenue
Cambridge, MA 02139
Tel: (617) 253-7277

Roger A. Duncan
Department Staff
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (702) 827-6244

Donald R. Edmonds
Department Staff
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6808

Timothy D. Evans
Computer Scientist
U.S. Army Research Office
P. O. Box 12211
Research Triangle Park, NC 27701
Tel: (919) 549-0641

Carl E. Fechner
Senior Scientist
Bolt Beranek and Newman, Inc.
10 Moulton Street
Cambridge, MA 02142
Tel: (617) 497-3335

Dorothy L. Finley
Team Leader
U.S. Army Research Institute for the
Behavior and Social Sciences
ATTN: PERI--SPB
5001 Eisenhower Avenue
Alexandria, VA 22333
Tel: (202) 274-8905

Joseph Frank
New Jersey Institute of Technology
323 High Street
Newark, NJ 07102

David R. Friedman
Member Technical Staff
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6418

Thomas D. Garvey
Program Director
SRI International
MS EK260
333 Ravenswood Avenue
Menlo Park, CA 94025
Tel: (415) 859-3486

Sherman Gee
Program Element Manager
Office of Naval Technology
(MAT 0733)
800 N. Quincy Street
Arlington, VA 22217
Tel: (301) 696-5120

Ralph Gerenz
Vice President
BETAC Corporation
76 Treble Cove Rd.
N. Billerica, MA 01862
Tel: (617) 272-6166

Paul E. Girard
CDR USN
Program Manager C² Technology
Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217
Tel: (202) 696-4713

Thomas W. Godowsky
Lt. USAF
Foreign Command and Control Systems
HQ Foreign Technology Division
Wright Patterson AFB, OH 45433
Tel: (513) 257-3556

Robert Goldsmith
Senior Member Technical Staff
BETAC Corporation
76 Treble Cove Road
N. Billerica, MA 01862
Tel: 617) 272-6166

Steven Goldstein
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6632

I. R. Goodman
Naval Ocean Systems Center
Code 7223, Seaside Bldg. 600, Rm 222
San Diego, CA 92152
Tel: (619) 225-2014

Jerren Gould
Senior Systems Engineer
Hughes Aircraft Company
P. O. Box 3310
Fullerton, CA 92634
Tel: (714) 732-1289

Conrad J. Grant
Senior Physicist
The Johns Hopkins University
Applied Physics Laboratory
Johns Hopkins Road
Laurel, MD 20707
Tel: (301) 953-7100, Ext. 3283

Frank Greitzer
Research Psychologist
Navy Personnel Research and Development
Center
Human Factors, Code 17-N
San Diego, CA 92152
Tel: (619) 225-2081

W. F. Griesse
Head, Software Systems Lab.
Hazeltine Corporation
Greenlawn, NY 11740
Tel: (516) 261-7000

Fred M. Griffiee
Tech. Advisor
Command Control and Communications Division
Code D10A
Marine Corps Development and Education Command
Quantico, VA 22134
Tel: (703) 640-2820

Bruce Hamill
Research Psychologist
Johns Hopkins University
Applied Physics Laboratory
Johns Hopkins Road
Laurel, MD 20707
Tel: (301) 953-7100, Ext. 572

James Henderson
Captain USN
Surface ASW Branch (OP-951D)
The PENTAGON
Room 50577
Washington, D.C. 20362
Tel: (202) 697-1066

Y. C. Ho
Professor EE
Harvard University
Pierce Hall
Cambridge, MA 02138
Tel: (617) 495-3992

Charles J. Holland
Group Leader Mathematician
Office of Naval Research
Code 411MA
800 N. Quincy Street
Arlington, VA 22217
Tel: (202) 696-4362

P. A. Humblet
Associate Professor EE
Massachusetts Institute of Technology
Room 35-203/LIDS
7. Massachusetts Avenue
Cambridge, MA 02139
Tel: (617) 253-7250

Leland T. Joe
Senior Engineer
M/A-COM Linkabit
1517 Westbranch Drive
McLean, VA 22102
Tel: (703) 734-8660

Randolph J. Johnstone
Engineer Scientist
Douglas Aircraft Company
3855 Lakewood Blvd.
Long Beach, CA 90846
Tel: 213) 593-8382

Glenn R. Linsenmayer
Fellow Engineer
Westinghouse Electric Corporation
Box 746-MS 434
Baltimore, MD
Tel: (301) 765-2243

P. Papantoni-Kazakos
Associate Professor
University of Connecticut
Box U-157
Storrs, CT 06268
Tel: (203) 486-2867

M. Craig Kerman
Program Manager
Lockheed Electronics Company
1501 US Highway 22
Plainfield, NJ 07061
Tel: (201) 757-1600, Ext. 2793

David Kleinman
Professor EECS
University of Connecticut
Box U-157
Storrs, CT 06268
Tel: (203) 486-3066

Rita E. Knox
Project Leader
Lockheed Electronics Company
1501 US Highway 22
CS#1, MS 252
Plainfield, NJ 07061
Tel: (201) 757-1600, Ext. 2464

Michael Kovacich
Scientist
Comptek Research Inc.
100 Corporate Place Suite 13
Vallejo, CA
Tel: (707) 552-3538

Leslie C. Kramer
Manager
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Thomas Kurien
Member Technical Staff
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Joel S. Lawson, Jr.
Technical Director
C-1 Systems and Technology Directorate
Naval Electronic Systems Command
Code 06T
Department of the Navy
Washington, D.C. 20363
Tel: (202) 692-9223

Kevin Laxar
Naval Submarine Medical Research Lab.
P. O. Box 900 SUBASE NLON
Groton, Connecticut 06349
Tel: (203) 449-2522

Alexander H. Levis
Senior Research Scientist
Massachusetts Institute of Technology
Room 35-410/LIDS
77 Massachusetts Avenue
Cambridge, MA 02139
Tel: (617) 253-7262

Robert Lobbia
ORINCON Corporation
3366 N. Torrey Pines Court
LaJolla, CA 92037
Tel: (619) 455-5530

Kenneth A. Loparo
Associate Professor SF
Case Western Reserve University
Cleveland, Ohio 44106
Tel: (216) 368-4115

Douglas Looze
Member Technical Staff
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Peter B. Luh
Assistant Professor EECS
University of Connecticut
Box U-157
Storrs, CT 06268
Tel: (203) 486-4821

John Machsdo
Command and Control Division
Naval Electronics Systems Command
Department of the Navy
Washington, D.C. 20360

Gerard S. Malecki
Office of Naval Research
Code 442
800 N. Quincy Street
Arlington, VA 22217
Tel: (202) 696-4741

James P. Marsh
President
Optimization Technology, Inc.
20380 Town Center Lane
Cupertino, CA 95014
Tel: (408) 973-1441

Donald Maurer
Mathematician
Johns Hopkins University
Applied Physics Laboratory
Johns Hopkins Road
Laurel, MD 20707
Tel: (703) 951-7100

John W. McCollum
Principal Analyst
Atlantic Analysis Corporation
Suite 219 Koger #5
Norfolk, VA 23502
Tel: (804) 461-1980

Raman K. Mehra
President
Scientific Systems, Inc.
54 Rindge Avenue Extension
Cambridge, MA 02140
Tel: (617) 661-6364

Michael Melich
Systems Engineer
Naval Research Laboratory
Code 7579
Washington, D.C. 20375
Tel: (202) 767-3959

Morton L. Metersky
Operations Research Analyst
Naval Air Development Center
Code 2031
Warminster, PA 18974
Tel: (215) 441-2378

William Meyers
R. Adm. USN, Ret.
9621 Ceralene Drive
Fairfax, VA 22032

Roger E. Miller
Research Specialist
Lockheed Missiles and Space
Box 17100 Orgn T1-50
Austin, TX 78760
Tel: (512) 448-5812

John R. Miller
TRW Defense Systems Group
One Space Park
MS 134/4085
Redondo Beach, CA 90278
Tel: (213) 217-4288

George Moeller
Director, Behavioral Sciences
Naval Submarine Medical research Lab
Box 900
Groton, CT 06349
Tel: (203) 449-3668

Paul H. Moose
Professor EE
Naval Postgraduate School
Monterey CA 93940
Tel: 408-646-2838

Peter D. Morgan
SCICON, Ltd.
49 Berners Street
London W1P 4AQ
ENGLAND
Tel: 01-218-4259

Franklin L. Moses
Acting Chief
Battlefield Information Systems Technical
US Army Research Institute
ATT: PERI-SF
5001 Eisenhower Avenue
Alexandria, VA 22333
Tel: (202) 274-8921

Alfred H. Mossler
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6000

Monte R. Mueller
Engineering
Boeing Aerospace
Slater Street South
Kirkland, WA 98033
Tel: (206) 773-1093

Carl N. Mutchler
The MITRE Corporation
Mail Stop W932
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6122

David Newport
Operations Research Analyst
Naval Weapons Center
ATTN: Code 3903
China Lake, CA 93555
Tel: (619) 939-2378

William O'Halloran
JAYCOR
300 Unicorn Park Drive
Woburn, MA 01801
Tel: (617) 933-6805

Ann T. Orlando
Senior Scientist
JAYCOR
300 Unicorn Park Drive
Woburn, MA 01801
Tel: (617) 933-6805

Glenn Osga
Human Factors Engineer
Code 8231
Naval Ocean Systems Center
San Diego, CA 92152
Tel: (619) 225-7372

Robert Overton
Comptek Research
2697 International Parkway
Virginia Beach, VA 23452
Tel: (804) 427-2505

Steven L. Feudergast
Sr. Staff Engineer
Hughes Aircraft Company
P. O. Box 3310
Building 691/B220
Fullerton, CA 92634
Tel: (714) 670-3264

Anastasios Perakis
Assistant Professor
NA/ME Building, North Campus
The University of Michigan
Ann Arbor, MI 48709
Tel: (313) 763-5006

Lloyd S. Peters
Associate Director
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
Tel: (415) 859-3650

Philip L. Phipps
Sperry Corporation
Computer Systems - MS CCA07
Sperry Park,
P. O. Box 43525
St. Paul, Minnesota 55164
Tel: (612) 456-6242

John G. Pierce
Associate Director
Ketrion, Inc.
1700 N. Moore Street
Arlington, VA 22209
Tel: (703) 558-8850

James R. Powell
LCdr. USN
SMC 1682
Naval Postgraduate School
Monterey, CA 93940
Tel: (408) 646-2056

William H. Rogers
Research Psychologist
Naval Submarine Medical Research Lab
P. O. Box 900
Naval Submarine Base New London
Groton, CT 06349
Tel: (203) 449-3668

Ol Rosenstark
New Jersey Institute of Technology
23 High Street
Newark, NJ 07102

Miss R. Sandell, Jr.
President
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

D. Roger Saxon
Statistician
Navy Personnel Research
and Development Center
Human Factors - Code 17N
San Diego, CA 92152
Tel: (619) 275-2081

Daniel Schutzer
Technical Director Naval Intelligence
Chief Naval Operations
Office of Naval Intelligence
(OP 009T)
Washington, D.C. 20350
Tel: (202) 697-2545

Brad D. Scott
Department of Psychology
University of Illinois
603 E. Daniel Street
Champaign, IL 61820
Tel: (217) 333-6195

Adrian Segal
Visiting Scientist
Thomas J. Watson Research Center
Yorktown Heights, NY 10598
Tel: (914) 945-3096

Robert A. Shade
Hazeltime Corporation
Cuba Hill Road
Greenlawn, NY 11740
Tel: (516) 261-7000 Ext. 2708

James G. Smith
Code 340R
Naval Air Systems Command
Washington, D.C. 20341
Tel: (202) 692-7414

Frank M. Snyder
Naval Operations
Naval War College
Newport, RI 02841
Tel: (401) 841-2396

David Spain
Manager
Advanced Information and Decision Systems
201 San Antonio Circle
Mountain View, CA 94040
Tel: (415) 941-3912

Debra A. Stabile
Sr. Associate
ICF Inc.
1850 K Street, Northwest
Suite 950
Washington, D.C. 20006
Tel: (202) 862-1100

Edward G. Stacy
Captain USN
The Pentagon
ASW Systems Project (RM-4)
Washington, D.C. 20362
Tel: (202) 692-1021

Fred Stamberger
Manager, System Engineering
and Analysis Department
Hughes Aircraft Company
1901 W. Malvern
Fullerton, CA 90632
Tel: (714) 732-1210

Peter J. E. Stan
Economist
The Rand Corporation
1700 Main Street
Santa Monica, CA 90406
Tel: (213) 393-0411

Stuart H. Starr
Director, Systems Planning
and Evaluation
M/A-COM Linkabit
1517 Westbranch Drive
McLean, VA 22102
Tel: (703) 734-8660

J. W. Stauffer
UEPO
GE Company
FRP 1-J7
Syracuse, NY 13210
Tel: (315) 456-4271

George Tasoulis
NAVELEX
NC #1 Code ELEX 611
Washington, D.C. 20363
Tel: (202) 692-6414

Demosthenis Teneketzis
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Robert R. Tenney
Asst. Professor EECS
35-213/LIDS
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
Tel: (617) 253-2359

Joel Trimble
Project Manager
Code 251
Office of Naval Research
Arlington, VA 22217
Tel: (202) 696-4713

Edison Tse
Professor
Dept. of Engineering-Economic Systems
Stanford University
Stanford, CA 94305
Tel: (415) 497-4777

Lena Valavani
Technical Staff
35-437A/LIDS
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
Tel: (617) 253-2157

Scott vanTouningen
Captain, USAF
Air Force Institute of Technology
1210 W. Washington Street
Champaign, IL 61821
Tel: (217) 351-5111

W. S. Vaughan
Leader
Engineering Psychology Group
Code 442FP
Office of Naval Research
Arlington, VA 22217
Tel: (202) 696-4290

Ronald Virostek
Group Leader
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-6266

Robert Washburn
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

John W. Watson
Member, Technical Staff
The MITRE Corporation
1820 Dolley Madison Blvd.
McLean, VA 22102
Tel: (703) 827-7355

Thomas M. Webb
General Engineer
Naval Surface Weapons Center
Code N51
Dahlgren, VA 22448
Tel: (703) 663-7431

Richard L. Weis
Member Technical Staff
Logicon, Inc.
21031 Ventura Blvd.
Suite 1200
Woodland Hills, CA 91364
Tel: (213) 887-4950

Moshe Weiss
J317 EQUAD
EECS Dept.
Princeton University
Princeton, NJ 08544
Tel: (609) 924-0497

Chelsea C. White, III
Professor
Department of Systems Engineering
University of Virginia
Thornton Hall
Charlottesville, VA 22901
Tel: (804) 924-5395

Joseph G. Wohl
Vice President
ALPHATECH, Inc.
2 Burlington Executive Park
111 Middlesex Turnpike
Burlington, MA 01803
Tel: (617) 273-3388

Stephen Wolff
Ballistic Research Laboratory
Aberdeen Proving Ground, MD 21005
Tel: (301) 278-6662

M. E. Womble
Staff Engineer, Sr.
Lockheed Missiles and Space Co., Inc.
2124 E St. Elmo Road
Austin, TX 78744
TEL: (512) 448-5875

Eugene Wong
Professor Dept. of EECS
University of California
Berkeley, CA 94720
TEL: (415) 642-5807

Raphael W. Wong
The MITRE Corporation
MS-W932
1820 Dolley Madison Blvd.
McLean, VA 22102
TEL: (703) 827-6122

Al M. Zeid
Director C³ Laboratory
Naval Postgraduate School
Code 74
Monterey, CA 93943
TEL: (408) 646-2772

STUDENT REGISTRATION

Erin Boettcher
35-401/LIDS
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-3444

Arthur Lee
35-401/LIDS
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-3444

Bo-Chu Chang
Dept. of EECS
University of Connecticut
P. O. Box U-157
Storrs, CT 06268
TEL: (203) 486-4691

Moses Ma
35-407/LIDS
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-3631

Li-Chung Chang
Dept. of EECS
University of Connecticut
P. O. Box U-157
Storrs, CT 06268
TEL: (203) 487-6089

Peter Ng
35-409/LIDS
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-5997

Jim Chann
Ashdown House #515B
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-9838

Jim Roseborough
Ashdown House #503-C
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 225-9838

Lioria H.L. Chyen
35-401/LIDS
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-3444

Daniel Serfaty
Dept. of EECS
University of Connecticut
P. O. Box U-157
Storrs, CT 06268
TEL: (203) 486-3261

Barry Ephraim
Ames Street
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 225-6416

Eric Soulsby
Dept. of EECS
University of Connecticut
P. O. Box U-157
Storrs, CT 06268
TEL: (203) 486-2210

Nikos Kapasouris
35-407/LIDS
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-3631

John Tsitaiklis
35-407/LIDS
Massachusetts Institute of Technology
Cambridge, MA 02139
TEL: (617) 253-3631

6th MIT/ONR WORKSHOP ON C³ SYSTEMS

JULY 25 TO JULY 29, 1983

FINAL PROGRAM
(AS OF JULY 22, 1983)

Sponsored by

LABORATORY FOR INFORMATION AND DECISION SYSTEMS
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
CAMBRIDGE, MASSACHUSETTS 02139

and

OFFICE OF NAVAL RESEARCH
DEPARTMENT OF THE NAVY
ARLINGTON, VIRGINIA 22217

Preceding page blank

MONDAY MORNING, JULY 25, 1983

SESSION 1: ROOM 10-250:

Chairman: Professor Michael Athans, MIT

8:30 - 9:00 A.M. REGISTRATION

9:00 - 9:30 A.M. INTRODUCTION AND WELCOME
Professor Michael Athans, MIT
Dr. Charles J. Holland, ONR

9:30 - 10:00 A.M. MULTI-SENSOR FUSION, COMMUNICATIONS AND INFORMATION WARFARE
D. Schutzer, Office of Naval Intelligence

10:00 - 10:30 A.M. C³ ASSESSMENT IN A MISSION AREA CONTEXT
L. T. Joe and S. H. Starr, M/A-COM Linkabit

10:30 - 11:00 A.M. BREAK

11:00 - 11:30 A.M. AUTOMATED WAR GAMING AS A TESTBED FOR EXAMINING ISSUES OF STRATEGIC C³ I
P. K. Davis and P. Stan, The RAND Corporation

11:30 - 12:00 P.M. LANCHESTER'S EQUATIONS AND MATRIX GAMES
J. M. Wozencraft and P. H. Moose
Naval Postgraduate School

12:00 - 1:30 P.M. LUNCH BREAK

MONDAY AFTERNOON, JULY 25, 1983

SESSION 2: ROOM 10-105: DETECTION AND LOCALIZATION

Chairman: Professor Robert R. Tenney, MIT

1:30 - 2:00 P.M. DETECTION MODELING OF UNSATURATED AND PARTIALLY SATURATED OCEAN ACOUSTIC SIGNALS
A. N. Perakis, University of Michigan
R. N. Psaraffis and H. I. Gonzales, MIT

2:00 - 2:30 P.M. FURTHER RESULTS IN MULTIPLATFORM CORRELATION AND INTERSHIP ALIGNMENT IN THE NAVAL BATTLEGROUP
M. Kovacich, Comptek Research, Corp.

2:30 - 3:00 P.M. THE USE OF RECIPROCAL POLAR CO-ORDINATES IN PASSIVE TRACKING
A. Bamford, SCICON, Ltd.

3:00 - 3:30 P.M. MANEUVERING FOR BEARINGS-ONLY ESTIMATION
G. Z. Wilhoit, U. S. Navy
R. R. Tenney, MIT

3:30 - 4:00 P.M. BREAK

4:00 - 4:30 P.M. GENERALIZED TRACKER/CLASSIFIER (GTC) - A SYSTEM FOR
TRACKING AND CLASSIFICATION OF MULTIPLE TARGETS

S. Mori, C. Y. Chong, and R. P. Wishner
Advanced Information and Decision Systems

E. Tse, Stanford University

4:30 - 5:00 P.M. OPTIMAL MANEUVER DETECTION AND ESTIMATION IN
MULTIOBJECT TRACKING

T. Kurien and A. Blitz, ALPHATECH, Inc.

5:00 - 5:30 P.M. A FRAMEWORK FOR SURVEILLANCE ALGORITHM EVALUATION

J. Weiss, Draper Laboratory
E. R. Tenney, MIT

MONDAY AFTERNOON, JULY 25, 1983

SESSION 3: ROOM 10-250: C³ SYSTEM EVALUATION

Chairman: Dr. Alexander H. Lewis, MIT

1:30 - 2:00 P.M. REQUIREMENTS FOR THE NAVY TACTICAL C³ SYSTEM

L. S. Peters, SRI International

2:00 - 2:30 P.M. A LARGE-SCALE CONTROL SYSTEM ANALYSIS METHODOLOGY

D. R. Friedman, The MITRE Corp.

2:30 - 3:00 P.M. TECHNIQUES FOR DETECTING COVER AND DECEPTION

R. Gorenz, BETAC, Corp.

3:00 - 3:30 P.M. A COMPARISON OF DEFENSE C³ MODELS

J.R. Dowdle, M.P. Merriman, R. F. Gendron, and
L.C. Kramer
ALPHATECH, Inc.

3:30 - 4:00 P.M. BREAK

4:00 - 4:30 P.M. OPTIMAL SHIP POSITIONS FOR NAVAL BATTLE GROUP DEFENSE

M. Athans, and R. C. Magonet-Meray, MIT

4:30 - 5:00 P.M. EVALUATING THE RESPONSE OF COMPLEX SYSTEMS TO
ENVIRONMENTAL THREATS: THE Σ T METHOD

G. C. Corynen, Lawrence Livermore National Laboratory

5:00 - 5:30 P.M. SOFTWARE FOR EXPLICITLY PROBABILISTIC MATHEMATICS

S. N. Goldstein, The MITRE Corporation

TUESDAY MORNING, JULY 26, 1983

SESSION 4: ROOM 10-250

Chairman: Professor Robert R. Tenney

8:30 - 9:00 A.M. REGISTRATION

9:00 - 9:30 A.M. SURVEILLANCE RESEARCH AT MIT
R. R. Tenney, MIT

9:30 - 10:00 A.M. DISTRIBUTED ESTIMATION IN THE MIT/LL DSN TEST-BED
J. R. Delaney, R. T. Lacoss, P. E. Green
Lincoln Laboratory

10:00 - 10:30 A.M. DECISIONMAKING: A COMMANDER'S CONSIDERATION
W. Meyers, U. S. Navy (Ret. Admiral)

10:30 - 11:00 A.M. BREAK

11:00 - 11:30 A.M. ARTILLERY CONTROL ENVIRONMENT
S. Wolff, Ballistic Research Laboratory

11:30 - 12:00 P.M. COMMAND DECISION MAKERS AND THEIR MODES OF INTERACTION
P. D. Morgan, SCICON, Ltd.

12:00 - 1:30 P.M. LUNCH BREAK

TUESDAY AFTERNOON, July 26, 1983

SESSION 5: ROOM 10-105: DISTRIBUTED ESTIMATION AND DATA FUSION

Chairman: Professor Robert R. Tenney, MIT

1:30 - 2:00 P.M. MULTI-SENSOR FUSION IN A MULTI-TARGET ENVIRONMENT
R. N. Lobbis and D. L. Alspach, ORINCON

2:00 - 2:30 P.M. DISTRIBUTED ESTIMATION SYSTEMS
C. Y. Chong, S. Mori, Advanced Information and
Decision Systems
E. Tse, Stanford University

2:30 - 3:00 P.M. DISTRIBUTED ESTIMATION WITH COMMUNICATION DELAY
A. Ozbek and R. R. Tenney, MIT

3:00 - 3:30 P.M. TRACK ASSOCIATION ALGORITHMS
A. Bamford, SCICON, Ltd.

3:30 - 4:00 P.M. BREAK

4:00 - 4:30 P.M. INCA: AN ENVIRONMENT FOR EXPLORATORY DEVELOPMENT OF
TACTICAL DATA FUSION TECHNIQUES
L. S. Gross and H. J. Payne, VERAC, Inc.

4:30 - 5:00 P.M. A UNIFIED APPROACH TO MODELING AND COMBINING OF
EVIDENCE THROUGH RANDOM SET THEORY
I. R. Goodman, NOSC

5:00 - 5:30 P.M. EXPERT SYSTEMS APPLIED TO AUTOMATED MILITARY
INTELLIGENCE
M. E. Womble, Lockheed Missile and Space

TUESDAY AFTERNOON, July 26, 1983

SESSION 6: ROOM 10-250 : ORGANIZATION MODELS AND DESIGN

Chairman: Dr. Alexander H. Lewis, MIT

1:30 - 2:00 P.M. DISTRIBUTED OPTIMIZATION ALGORITHMS WITH
COMMUNICATIONS
J. N. Tsitsiklis and M. Athans, MIT

2:00 - 2:30 P.M. A LAYERED MODEL FOR COMMAND-CONTROL
W. F. Griesse, Hazeltine Corp.

2:30 - 3:00 P.M. MODELING THE ASW TACTICAL C³ DECISION PROCESS
M. Alexandridis, J. Deckert, E. Ratin, and J. G. Wohl
ALPHATECH, Inc.

3:00 - 3:30 P.M. THE DESIGN OF INFORMATION STRUCTURES:
BASIC ALLOCATION STRATEGIES FOR ORGANIZATIONS
D. A. Stabile, ICF, Inc.
A. H. Lewis, MIT

3:30 - 4:00 P.M. BREAK

4:00 - 4:30 P.M. MESSAGE STANDARDS IN SOFTWARE ENGINEERING
J. V. Bronson, Col. U. S. Marine Corps

4:30 - 5:00 P.M. MULTI-OBJECTIVE DECISIONMAKING, INCOMPLETE INFORMATION
AND COORDINATION PROBLEM IN C³ SYSTEMS
K. Loparo and M. Mesarovic
Case Western Reserve University

WEDNESDAY MORNING, JULY 26, 1983

SESSION 7: ROOM 10-25

Chairman: Professor Robert R. Tenney, MIT

8:30 - 9:00 A.M.	REGISTRATION
9:00 - 9:30 A.M.	COMMAND AND CONTROL ORGANIZATIONS FOR NAVAL BATTLE FORCES M. Athapa, MIT
9:30 - 10:00 A.M.	INDIVIDUAL DIFFERENCES IN MILITARY DECISIONMAKING: THE CLASSIFICATION PERFORMANCE OF ACTIVE SONAR OPERATORS J. G. Wohl, ALPHATECH, Inc.
10:00 - 10:30 A.M.	DECISION AND DISPLAY ANALYSIS IN A SIMPLE SURVEILLANCE PROBLEM R. L. Hershman, F. L. Greitzer, NPRDC
10:30 - 11:00 A.M.	<u>BREAK</u>
11:00 - 11:30 A.M.	DISTRIBUTED ASYNCHRONOUS ALGORITHMS D. P. Bertsekas, MIT
11:30 - 12:00 P.M.	HYBRID ROUTING IN COMMUNICATION NETWORKS A. Segall, J. M. Jaffe and F. H. Moss IBM Thomas J. Watson Research Center
12:00 - 1:30 P.M.	<u>LUNCH BREAK</u>

WEDNESDAY AFTERNOON, JULY 27, 1983

SESSION 8: ROOM 10-105: COMMUNICATIONS

Chairman: Professor Robert R. Tenney, MIT

1:30 - 2:00 P.M.	INTERSHIP SENSOR ALIGNMENT USING UHF TACTICAL DATA LINK RELATIVE NAVIGATION R. H. Overton, Comptek Research, Inc.
2:00 - 2:30 P.M.	COMMUNICATIONS SUPPORT OF DISTRIBUTED COMMAND AND CONTROL G. A. Clapp, NOSC
2:30 - 3:00 P.M.	THE USE OF SUMMARY DISPLAY TRANSFERS FOR COORDINATION BETWEEN COMMAND FACILITIES FOR THE COMPOSITE WARFARE COMMANDER AND THE ANTI-AIR WARFARE COMMANDER OF A NAVY BATTLE GROUP C. J. Grant, The Johns Hopkins University

3:00 - 3:30 P.M. BREAK

3:30 - 4:00 P.M. DATA BASE REDUNDANCY IN VULNERABLE COMMUNICATION NETWORKS
M. Ma, and M. Athans, MIT

4:00 - 4:30 P.M. A COLLISION RESOLUTION PROTOCOL WITH LIMITED CHANNEL SENSING FINITELY MANY USERS
P. Papantoni-Kazakos, G. D. Marcus, M. Georgiopoulos
University of Connecticut

4:30 - 5:00 P.M. OPEN

5:00 - 6:00 P.M. WINE AND CHEESE (BUSH ROOM 10-105)

WEDNESDAY AFTERNOON, JULY 27, 1983

SESSION 9: ROOM 10-250: HUMAN DECISION MODELS

Chairman: Dr. Alexander H. Levis, MIT

1:30 - 2:00 P.M. HUMAN DECISIONMAKING IN DYNAMIC ENVIRONMENT WITH INCREASING INFORMATION PROCESSING DEMANDS
D. Serfaty, University of Connecticut

2:00 - 2:30 P.M. C³I SYSTEM DESIGN AND DEVELOPMENT: INSIGHT FROM THE BEHAVIORAL SCIENCE PERSPECTIVE
D. L. Finley, Army Research Institute
W. P. Cherry, Vector Research, Inc.

2:30 - 3:00 P.M. MODELS OF INFORMATION STORAGE AND MEMORY IN C³ SYSTEMS
S. A. Hall, FCA, Inc.
A. H. Levis, MIT

3:00 - 3:30 P.M. BREAK

3:30 - 4:00 P.M. THE COGNITIVE ORGANIZATION OF SUBMARINE SONAR INFORMATION: A MULTIDIMENSIONAL SCALING ANALYSIS
V. Laxar, G. Moeller, and W. Rogers
Naval Submarine Medical Research Laboratory

4:00 - 4:30 P.M. THE EFFECTS OF SPATIAL INFORMATION DISPLAYS ON DECISION MAKING PERFORMANCE IN TACTICAL C³ SYSTEMS
B. D. Scott and C. D. Wickens
University of Illinois

4:30 - 5:00 P.M. DEVELOPMENT OF A GENERALIZED HUMAN-MACHINE INTERFACE
R. E. Knox, Lockheed Electronics Co., Inc.

5:00 - 6:00 P.M. WINE AND CHEESE (Bush Room 10-105)

THURSDAY MORNING, JULY 28, 1983

SESSION 10: ROOM 10-250

Chairman: Dr. Alexander H. Lewis, MIT

8:30 - 9:00 A.M. REGISTRATION

9:00 - 9:30 A.M. PROBLEMS IN THE ANALYSIS AND EVALUATION OF
INFORMATION PROCESSING AND DECISIONMAKING
OR NIZATIONS
A. H. Lewis, MIT

9:30 - 10:00 A.M. DOING C² EXPERIMENTS WITH WAR GAMES
J. S. Lawson, Jr., NAVELEX

10:00 - 10:30 A.M. DESIGN OF A SYSTEM EXECUTIVE FOR THE MANAGEMENT OF
COMMAND/CONTROL FACILITIES
E. R. Ducot, MIT

10:30 - 11:00 A.M. BREAK

11:00 - 11:30 A.M. A KNOWLEDGE BASED, INTERACTIVE PROCEDURE FOR PLANNING
AND DECISION SUPPORT UNDER UNCERTAINTY AND PARAMETER
IMPRECISION
A. P. Sage and C. C. White, III
University of Virginia

11:30 - 12:00 P.M. DATA INTEGRATION FOR COMMAND AND CONTROL SUPPORT
SYSTEMS
J. C. Machado, NAVELEX

12:00 - 1:30 P.M. LUNCH BREAK

THURSDAY AFTERNOON, JULY 28, 1983

SESSION 11: ROOM 10-105: THEORETICAL ADVANCES

Chairman: Professor Robert R. Tenney, MIT

1:30 - 2:00 P.M. A MINIMUM SENSITIVITY INCENTIVE CONTROL APPROACH TO
TEAM PROBLEMS

T. Basar, J. B. Cruz, Jr., and D. Cansever
University of Illinois

2:00 - 2:30 P.M. SUBJECTIVE GAMES OF INCOMPLETE INFORMATION

D. Teroketzi and D. Castanon, ALPHATECH, Inc.

2:30 - 3:00 P.M. OPTIMAL STRUCTURES IN DYNAMIC TEAM PROBLEMS

R. R. Tenney, MIT

3:00 - 3:30 P.M. ON THE COMPLEXITY OF DISTRIBUTED DECISION PROBLEMS
J. Tsitsiklis and M. Athans, MIT

3:30 - 4:00 P.M. BREAK

4:00 - 4:30 P.M. ON THE ASSIGNMENT OF LEADERSHIP IN STACKELBERG GAMES
K. Loparo, Y. Choe, and S. Kahne
Case Western Reserve University

4:30 - 5:00 P.M. DYNAMIC, HIERARCHICAL DECISION PROBLEMS
P. B. Luh, and T. S. Chang, University of Connecticut

5:00 - 5:30 P.M. ON THE CONTROL OF SYSTEMS WITH ABRUPTLY CHANGING
STRUCTURE
B. E. Griffiths and K. A. Loparo
Case Western Reserve University

THURSDAY AFTERNOON, JULY 28, 1983

SESSION 12: ROOM 10-250: DECISION AIDS

Chairman: Ms. Elizabeth R. Ducot, MIT

1:30 - 2:00 P.M. A MAN-MACHINE INTERFACE CONCEPT FOR A STATE-OF-ART,
PROGRAMMABLE, SHIPBOARD, COMMAND/CONTROL CONSOLE
G. Osga and R. Fleming, NOSC

2:00 - 2:30 P.M. TACTICAL DECISION AIDS WHICH QUANTIFY JUDGEMENT
V. Brown, Decision Science Consortium, Inc.

2:30 - 3:00 P.M. EVIDENTIAL REASONING FOR SITUATION ASSESSMENT
T. D. Garvey and J. D. Lowrance, SRI International

3:00 - 3:30 P.M. AN EVALUATION OF ARIADNE
C. C. White, III and A. P. Sage
University of Virginia

3:30 - 4:00 P.M. BREAK

4:00 - 4:30 P.M. PERSONALIZED AND DESCRIPTIVE DECISION AIDS
M. S. Cohen, MAXIMA Corp.

4:30 - 5:00 P.M. SIM: A SMART INTERFACE TO A MODEL
A. L. Blitz, R. R. Tenney, L. C. Kramer
ALPHATECH, Inc.

5:00 - 5:30 P.M. ON THE NEED FOR WELL STRUCTURED C² EXPERIMENTS
F. C. Deckelman, NAVELEX

AUTHOR INDEX

	<u>Page</u>		<u>Page</u>
Athans, M.	215	Ma, M. H.	255
	232	Marcus, G. D.	268
	255	Merriman, M. P.	18
Bamford, A. C.	188	Moeller, G.	87
	198	Moose, P. H.	24
Basar, T.	243	Morgan, P. D.	61
Beagle, E.M.L.	188	Mori, S.	158
	198		172
Bennett, W.	10	Ning, Y.	249
Bertsekas, D. P.	39		
	215	Osga, G. A.	133
Blitz, A.	164	Overton, R. H.	195
Brown, R. V.	106		
		Patel, S.	198
Cansever, D.	243	Payne, H. J.	183
Castanon, D.A.	226	Perakis, A. N.	207
Chang, Y-S	249	Psaraffis, H. N.	207
Chong, Y. C.	158		
	172	Rogers, W. H.	87
Cohen, M.S.	106		
	129	Sage, A. P.	109
Cruz, Jr., J. B.	243		117
		Scherer, W. T.	117
Davis, P. K.	10	Schutzer, D.	139
Delaney, J. R.	152	Scott, B.	96
Dowdle, J. R.	18	Serfaty, D.	76
Dozono, S.	117	Soulsby, E. P.	76
		Stabile, D. A.	219
East, D.	188	Stan, J.E.	10
Gendron, R. F.	18	Tencketzis, D.	226
Georgopoulos, M.	268	Tenney, R. R.	178
Gerenz, R. F.	145		202
Goldsmith, R. P.	145		238
Goldstein, S. N.	48	Tse, E.	158
Gonzales, H. I.	207		172
Goodman, I. R.	42	Tsitsiklis, J.	215
Grant, C. J.	129		232
Green, P. E.	152		
Greitzer, F. L.	91	Washburn, R.	164
Gross, S. L.	183	Weiss, J. L.	178
		White, III, C.C.	109
Hall, S. A.	67		117
Hershman, R. L.	91	Wickens, C. D.	96
		Wilhoit, G. Z.	202
Kazakos-Papantoni, P	268	Willsky, A.	164
Kleinman, D. L.	76	Wishner, R. P.	172
Knox, R. E.	100	Wohl, J. G.	83
Kovacich, M.A.	191	Wozencraft, J. M.	24
Kramer, L. C.	18		
Kurien, T.	164		
Lacoss, R. T.	152		
Lawson, Jr., J. S.	3		
Lazar, K.	87		
Lee, J.	198		
Levin, A. M.	30		
	67		
	219		
Luh, P. B.	249		

Preceding page blank